



Maxtor Atlas 10K V

Product Manual

September 6, 2005
Revision 3

PN: 00001911

Maxtor® March 24,2005 Maxtor Corporation. All rights reserved. Printed in U.S.A.

This publication could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein – which will be incorporated in revised editions of the publication. Maxtor may make changes or improvements in the product(s) described in this publication at any time and without notice.

UL/CSA/VDE/TUV

UL standard 1950 3rd Edition recognition granted under File No. E146611

CSA standard C22.2-950 recognized under UL File No. E146611

TUV Rheinland EN 60950 Bauart Certificate R 72030235

Tested to FCC Rules for Radiated and Conducted Emissions, Part 15, Sub Part J, for Class-B Equipment.

Atlas 10K V model number 8JXXXXX meets the EU directive for the Restriction and Use of Hazardous Substances (RoHS), 2002/95/EC of the European Parliament and the council of 27 January, 2003. Atlas 10K V model number 8DXXXXX do not meet these initiatives.

PATENTS

These products are covered by or licensed under one or more of the following U.S. Patents:

4,419,701; 4, 538,193 4,625,109; 4,639,798; 4,647,769; 4,647,997; 4,661,696; 4,669,004; 4,675,652; 4,703,176; 4,730,321; 4,772,974; 4,783,705; 4,819,153; 4,882,671; 4,920,442; 4,920,434; 4,982,296; 5,005,089; 5,027,241; 5,031,061; 5,084,791; 5,119,254; 5,160,865; 5,170,229; 5,177,771; Other U.S. and Foreign Patents Pending.

Maxtor®, MaxFax® and No Quibble Service® are registered trademarks of Maxtor Corporation, registered in the U.S.A. and other countries. Maxtor Atlas, DisCache, WriteCache are trademarks of Maxtor Corporation. All other brand names or trademarks are the property of their manufacturers.

Maxtor reserves the right to make changes and improvements to its products, without incurring any obligation to incorporate such changes or improvements into units previously sold or shipped.

This product or document is protected by copyright and distributed under licences restricting its use, copying, distributing, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Maxtor and its licensors, if any.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraphs (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 and FAR 52.227-19.

THIS PUBLICATION IS PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.

You can request Maxtor publications from your Maxtor Sales Representative or order them directly from Maxtor.

Publication Number: PN: 000001911

Corporate Headquarters

500 McCarthy Blvd
Milpitas, California 95035
Tel: 408-894-5000
Fax: 408-362-4740

Research and Development Center

333 South Street
Shrewsbury, MA 01545
Tel: 1-800-2-MAXTOR or 1-800-262-9867

Before You Begin

Thank you for your interest in Maxtor hard disk drives. This manual provides technical information for OEM engineers and systems integrators regarding the installation and use of Maxtor hard drives. Drive repair should be performed only at an authorized repair center. For repair information, contact the Maxtor Customer Service Center at 800-2MAXTOR or 303-678-2045.

CAUTION: Maxtor hard drives are precision products. Failure to follow these precautions and guidelines outlined here may lead to product failure, damage and invalidation of all warranties.

- 1 BEFORE** unpacking or handling a drive, take all proper electro-static discharge (ESD) precautions, including personnel and equipment grounding. Stand-alone drives are sensitive to ESD damage.
- 2 BEFORE** removing drives from their packing material, allow them to reach room temperature.
- 3** During handling, *NEVER* drop, jar, or bump a drive.
- 4** Once a drive is removed from the Maxtor shipping container, *IMMEDIATELY* secure the drive through its mounting holes within a chassis. Otherwise, store the drive on a padded, grounded, antistatic surface.
- 5 NEVER** switch DC power onto the drive by plugging an electrically live DC source cable into the drive's connector. *NEVER* connect a live bus to the drive's interface connector. (P68 only)

Please do not remove or cover up Maxtor factory-installed drive labels. They contain information required should the drive ever need repair.

Regulatory Statements

This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to Part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. Any modifications of this device – unless expressly approved by the manufacturer – can void the user's authority to operate this equipment under Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) This device may not cause harmful interference and (2) This device must accept any interference that may cause undesirable operation.

This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult the dealer or an experienced radio/TV technician for help.

NOTE: *Additional information on the need to interconnect the device with shielded (data) cables or the need for special devices, such as ferrite beads on cables, is required if such means of interference suppression was used in the qualification test for the device. This information will vary from device to device and needs to be obtained from the EMC group or product manager.*

This Class B digital apparatus complies with Canadian ICES-003.

Table of Contents

Chapter 1

ABOUT THIS MANUAL

1.1	AUDIENCE	1-1
1.2	MANUAL ORGANIZATION	1-1
1.3	TERMINOLOGY AND CONVENTIONS	1-2
1.4	REFERENCES	1-3

Chapter 2

GENERAL DESCRIPTION

2.1	PRODUCT OVERVIEW	2-1
2.2	KEY FEATURES	2-1
2.3	REGULATORY COMPLIANCE STANDARDS	2-3
2.4	HARDWARE REQUIREMENTS	2-4

Chapter 3

INSTALLATION

3.1	SAFETY, HANDLING, & ELECTROSTATIC DISCHARGE PROTECTION	3-1
3.1.1	Safety Precautions	3-1
3.1.2	Handling	3-1
3.1.3	Electrostatic Discharge (ESD) Protection	3-2
3.2	SPACE REQUIREMENTS	3-4
3.2.1	Shock Feet	3-5
3.3	UNPACKING INSTRUCTIONS	3-6
3.4	HARDWARE OPTIONS	3-8
3.4.1	Configuration Jumpers and Connections	3-8
3.4.2	Jumper Options on the 68-Pin Wide PCB	3-8
3.4.2.1	Write Protection	3-10
3.4.2.2	Delay Spin (DS) Stagger SPin (SS)	3-11
3.4.2.3	Single-Ended Operation _ Force SE)Disable LVD).....	3-11
3.4.2.4	Remote Busy and Fault Displays.....	3-11
3.4.3	SCA-2 80-Pin Connector Versions	3-12
3.4.3.1	Termination for the 80-Pin SCA-2 Connector Versions.....	3-12
3.4.3.2	SCSI ID for SCA-2 Versions.....	3-13
3.4.3.3	Spin Up for the 80 Pin SCA-2 Connector Versions.....	3-14
3.4.3.4	Activity LED for SCA Connector Versions.....	3-15
3.5	INTERFACE CONNECTOR (J1).....	3-15
3.5.1	68-Pin Wide SCSI Connector - LVD	3-18
3.5.1.1	68-Pin Wide Connector.....	3-18
3.5.2	80-Pin SCA-2 SCSI Connector - LVD	3-19

3.5.2.1 80 Pin SCA-2 Mating Connectors.....	3-20
3.5.3 68-Pin Wide Single-Ended SCSI Connector	3-21
3.5.4 80-Pin SCA-2 Single-Ended SCSI Connector	3-22
3.6 DRIVE MOUNTING and installation	3-26
3.6.1 Orientation	3-26
3.6.2 Mounting Screw Clearance	3-27
3.6.3 Mounting	3-28
3.6.4 Ventilation	3-28

Chapter 4

SPECIFICATIONS

4.1 SPECIFICATION SUMMAR	4-1,4-2
4.2 FORMATTED CAPACITY	4-3
4.3 DATA TRANSFER RATES	4-4
4.4 TIMING SPECIFICATIONS	4-4
4.5 POWER.....	4-5
4.5.1 Power Sequencing	4-5
4.5.2 Power Reset Limits	4-5
4.5.3 Drive Power Dissipation	4-6
4.6 ACOUSTICS.....	4-7
4.7 MECHANICAL	4-8
4.8 ENVIRONMENTAL CONDITIONS.....	4-8
4.9 ELECTROMAGNETIC CONDITIONS	4-9
4.9.1 EMI/RFI Susceptibility	4-9
4.9.2 ESD	4-9
4.9.3 Sensitivity to Magnetic Fields	4-9
4.10 SHOCK AND VIBRATION.....	4-9
4.11 RELIABILITY	4-10
4.12 DISK ERRORS.....	4-10

Chapter 5

SCSI Description

5.1 Overview of the SCSI Command Descriptions.....	5-1
5.2 Command Descriptor Block.....	5-4
5.3 Status/Error Reporting.....	5-7
5.4 Auto Contingent Allegiance Condition and Contingent Allegiance Condition	5-8
5.5 Extended Contingent Allegiance Condition	5-9
5.6 Linked Commands	5-10
5.7 DATA Transfer Command Components.....	5-11
5.8 SCSI COMMAND DESCRIPTIONS	5-13
5.9 NEW OPERATING DEFINITION (40h)	5-14
5.10 FORMAT UNIT Command (04h).....	5-15
5.10.1 Five Forms of FORMAT UNIT Commands	5-17
5.10.2 FORMAT UNIT Parameter List	5-17
5.11 INQUIRY Command (12h).....	5-21
5.11.1 Standard Inquiry Data Page	5-22

5.11.2 Vital Product Data Pages	5-24
5.12 LOG SELECT Command (4Ch).....	5-32
5.13 LOG SENSE Command (4Dh).....	5-33
5.13.1 LOG SENSE Command Descriptor Block	5-34
5.13.2 LOG SENSE Log Pages	5-35
5.14 MODE SELECT (6) Command (15h).....	5-37
5.14.1 Initiator-Changeable Mode Pages	5-37
5.14.2 Mode Page Types	5-39
5.14.3 Mode Parameter List	5-39
5.14.4 Categories of Changeable Pages	5-41
5.14.5 Unit Attention Condition Page (00h)	5-41
5.14.6 Read-Write Error Recovery Page (01h)	5-42
5.14.7 Disconnect-Reconnect Page (02h)	5-44
5.14.8 Verify Error Recovery Page (07h)	5-45
5.14.9 Caching Page (08h)	5-46
5.14.10 Control Mode Page (0Ah)	5-48
5.14.11 Notch and Partition Page (0Ch)	5-50
5.14.12 Port Control Mode Page (19h)	5-52
5.14.14 Information Exceptions Control Page (1Ch)	5-59
5.15 MODE SELECT (10) Command (55h).....	5-62
5.16 MODE SENSE (6) Command (1Ah).....	5-64
5.16.1 Read-Only Mode Pages	5-65
5.16.2 Format Device Page (03h)	5-65
5.16.3 Rigid Disk Geometry Page (04h)	5-67
5.17 MODE SENSE (10) Command (5Ah).....	5-69
5.18 PERSISTENT RESERVATION IN Command (5Eh).....	5-71
5.19 PERSISTENT RESERVATION OUT Command (5Fh).....	5-79
5.20 READ (6) Command (08h).....	5-89
5.21 READ (10) Command (28h).....	5-90
5.22 READ BUFFER Command (3Ch).....	5-91
5.23 READ CAPACITY Command (25h).....	5-93
5.24 READ DEFECT DATA Command (10) (37h).....	5-94
5.25 READ DEFECT DATA Command (12) (B7h).....	5-97
5.26 READ LONG Command (3Eh).....	5-99
5.27 REASSIGN BLOCKS Command (07h).....	5-100
5.28 RECEIVE DIAGNOSTIC RESULTS Command (1Ch).....	5-101
5.28.1 Supported Diagnostics Pages Page (00h)	5-102
5.28.2 Translate Address Page (40h)	5-102
5.29 RELEASE (6) Command (17h).....	5-104
5.30 RELEASE (10) Command (57h).....	5-105
5.31 REPORT DEVICE IDENTIFIER Command (A3h).....	5-106
5.32 REPORT LUNS Command (A0h)	5-107
5.33 REQUEST SENSE Command (03h).....	5-108
5.33.1 Sense Data Availability	5-108
5.33.2 Status Reporting	5-109
5.33.3 Sense Data Format for Error Code 70h and Error Code 71h	5-109

5.34 RESERVE (6) Command (16h)	5-124
5.35 RESERVE (10) Command (56h)	5-125
5.36 REZERO UNIT Command (01h)	5-127
5.37 SEEK (6) Command (0Bh)	5-128
5.38 SEEK (10) Command (2Bh)	5-129
5.39 SEND DIAGNOSTIC Command (1Dh)	5-130
5.39.1 Supported Diagnostic Page List	5-131
5.39.2 Translate Address Page	5-131
5.40 SET DEVICE IDENTIFIER Command (A4h)	5-132
5.41 START STOP UNIT Command (1Bh)	5-133
5.42 SYNCHRONIZE CACHE Command (35h)	5-134
5.43 TEST UNIT READY Command (00h)	5-135
5.44 VERIFY Command (2Fh)	5-136
5.45 WRITE (6) Command (0Ah)	5-137
5.46 WRITE (10) Command (2Ah)	5-138
5.47 WRITE AND VERIFY Command (2Eh)	5-139
5.48 WRITE BUFFER Command (3Bh)	5-140
5.49 WRITE LONG Command (3Fh)	5-142

Chapter 6

FEATURE DESCRIPTIONS

6.1 KEY FEATURES	6-1,6-2
6.2 WRITE-BACK CACHING	6-2,6-3
6.3 ZERO LATENCY READ/WRITE	6-3
6.4 DISCONNECT-RECONNECT	6-3
6.5 TRACK AND CYLINDER SKEWING	6-3
6.6 AVERAGE ACCESS TIME	6-3
6.7 EMBEDDED SERVO SYSTEM	6-3,6-4
6.8 DATA INTEGRITY AND SECURITY	6-4
6.8.1 Media Error Protection	6-4
6.8.2 Transfer Error Protection	6-4
6.8.3 Addressing Error Protection	6-4
6.8.4 Data Sector Reallocation Error Protection	6-4
6.8.5 Data Verification	6-5
6.9 TAGGED COMMAND QUEUING	6-5
6.10 COMMAND REORDERING	6-5
6.11 BANDED RECORDING	6-5
6.12 SPECIAL FUNCTIONS	6-5
6.13 DIAGNOSTICS	6-6
6.13.1 Power On Self Test (POST)	6-6
6.13.2 Periodic Self-Adjustments	6-6
6.13.3 Host Diagnostics	6-7
6.13.4 Log Sense	6-7
6.14 ERROR RECOVERY	6-7
6.14.1 Seek Error Recovery	6-7

6.14.2 Data Read Errors	6-7
6.14.3 Reallocation of Bad Blocks	6-8
6.15 THE ULTRA160 and ULTRA320 LOW VOLTAGE DIFFERENTIAL (LVD)	
SCSI INTERFACE	6-8,6-9
6.15.1 DIFFSENS	6-10
6.15.2 DIFFSENS SWITCHING	6-10
6.15.3 DT AND ST CLOCKING	6-11
6.15.4 CYCLIC REDUNDANCY CHECKING	6-12
6.15.5 DOMAIN VALIDATION	6-12
6.15.6 Adaptive Active Filter (AAF)	6-12
6.15.7 Transmitter Pre-Compensation with Cutback	6-13
6.15.8 Free Running Clock (FRC)	6-13
6.15.9 Skew Compensation	6-13
6.15.10 Information Unit Transfers	6-13
6.16 HOT PLUGGING/REMOVAL AND INSERTION OF SCSI DEVICES	6-13
6.16.1 System Considerations	6-14
6.16.2 APPLICABLE SCSI PHYSICAL DOCUMENTS	6-15

Appendix A

QUICK REFERENCE

A.1 SCSI-2/SCSI-3 Equivalent Terminology	A-1
A.2 SCSI Commands and Messages.....	A-2
A.3 SCSI Pages	A-3
A.4 Sense Keys	A-5
A.5 Status Codes.....	A-5

Appendix B

NEGOTIATED RATE INFORMATION PAGE REFERENCE

B.1 TRANSFER PERIOD FACTOR FIELD.....	B-1
B.2 Settings for the QAS, DT, and IU fields	B-2

List of Figures

Figure 3-1	Dimensions for the Maxtor Atlas 10K V Hard Disk Drives	3-3
Figure 3-2	Drive Packing Assembly	3-5
Figure 3-3	Jumper Locations on the 68-Pin Wide SCSI Drive PCB	3-7
Figure 3-4	Pin Locations on SCA-2 Connector	3-11
Figure 3-5	J1 Interface Connector Configurations.....	3-15
Figure 3-6	Mounting Dimensions Maxtor Atlas 10K V Drives.....	3-23
Figure 3-7	Mounting Screw Clearance for Maxtor Atlas 10K V Drive	3-24
Figure 3-8	Lengthwise Airflow Cooling.....	3-25
Figure 3-9	Crosswise Airflow Cooling.....	3-26
Figure 6-1	ST and DT CLocking	6-11

List of Tables

Table 3-1	SCSI ID Selection on Option Connector (68-Pin SCSI Connector Drives)	3-8
Table 3-2	SCSI ID Pin Assignments (SCA-2 Connector Versions of the Disk Drive)	3-12
Table 3-3	Spin Up on Power On Options	3-13
Table 3-4	68-Pin Wide LVD Pin Assignments	3-16
Table 3-5	80-Pin SCA-2 LVD Pin Assignments	3-17
Table 3-6	68-Pin Wide Single-Ended Pin Assignments	3-19
Table 3-7	80-Pin SCA-2 Single-Ended Pin Assignments	3-20
Table 4-1	Specifications	4-1
Table 4-2	Formatted Capacity	4-3
Table 4-3	Timing Specifications for Maxtor Atlas 10K V	4-4
Table 4-4	Power Reset Limits ¹	4-5
Table 4-5	Power Dissipation in Various Modes (Low-Profile Drives)	4-6
Table 4-6	Acoustical Characteristics—Sound Power per ISO 7779	4-7
Table 4-7	Environmental Specifications	4-8
Table 4-8	Shock and Vibration Specifications	4-10
Table 4-9	Error Rates	4-11
Table 5-1	Supported SCSI Commands	5-1
Table 5-2	Typical (6-Byte) Command Descriptor Block-Data Format	5-4
Table 5-3	Command Descriptor Block-Field Descriptions	5-4
Table 5-4	Typical (6-Byte) Command Descriptor Block - Data Format	5-5
Table 5-5	Command Descriptor Block Control Field-Field Descriptions	5-6
Table 5-6	Status Codes	5-7
Table 5-7	Length Fields	5-11
Table 5-8	DATA-Phase Command Contents	5-11
Table 5-9	CHANGE DEFINITION Command Descriptor Block-Data Format	5-14
Table 5-10	CHANGE DEFINITION-Field Descriptions	5-14
Table 5-11	FORMAT UNIT Command Descriptor Block-Data Format	5-16
Table 5-12	FORMAT UNIT Command-Field Descriptions	5-16
Table 5-13	FORMAT UNIT Command Supported Options	5-17
Table 5-14	FORMAT UNIT Parameter List-Data Format	5-17
Table 5-15	.FORMAT UNIT Defect List Header-Data Format	5-18
Table 5-16	FORMAT UNIT Defect List Header-Field Descriptions	5-18
Table 5-17	FORMAT UNIT Defect Descriptor-Block Format	5-19
Table 5-18	FORMAT UNIT Defect Descriptor- Physical Sector and Bytes From Index Format	5-19
Table 5-19	FORMAT UNIT Initialization Pattern Descriptor-Data Format	5-19
Table 5-20	FORMAT UNIT Initialization Pattern Descriptor-Field Descriptions	5-20
Table 5-21	FORMAT UNIT Initialization Pattern Type	5-20

Table 5-22	INQUIRY Command Descriptor Block-Data Format	5-21
Table 5-23	INQUIRY Command Descriptor Block-Field Descriptions	5-21
Table 5-24	Standard Inquiry Data Page-Data Format.....	5-22
Table 5-25	Standard Inquiry Data Page-Field Descriptions	5-23
Table 5-26	Supported Vital Product Data Pages -Data Format.....	5-24
Table 5-27	Vital Product Data-Page Codes	5-24
Table 5-28	Unit Serial Number Page-Data Format.....	5-24
Table 5-29	Unit Serial Number Page-Field Descriptions	5-24
Table 5-30	Implemented Operating Definition Page-Data Format.....	5-25
Table 5-31	Implemented Operating Definition Page-Field Descriptions	5-25
Table 5-32	ASCII Implemented Operating Definition Page - Data Format.....	5-26
Table 5-33	Device Identification Page-Data Format	5-26
Table 5-34	Device Identification Page-Field Description	5-27
Table 5-37	Command Support Data Page-Data Format.....	5-28
Table 5-38	Command Support Data Page-Field Descriptions	5-29
Table 5-39	Command Support Data Page Command or Operation Codes	5-30
Table 5-40	LOG SELECT Command Descriptor Block-Data Format.....	5-32
Table 5-41	LOG SELECT Command Descriptor Block-Field Descriptions.....	5-32
Table 5-42	Disk Drive Log Pages	5-33
Table 5-43	LOG SENSE Command Descriptor Block-Data Format.....	5-34
Table 5-44	LOG SENSE Command Descriptor Block-Field Descriptions	5-34
Table 5-45	LOG SENSE Log Page Format-Data Format.....	5-35
Table 5-46	LOG SENSE Log Page Format-Field Descriptions	5-35
Table 5-47	Generic Log Parameter-Data Format	5-35
Table 5-48	Generic Log Parameter-Field Descriptions.....	5-36
Table 5-49	MODE SELECT (6) Command Descriptor Block-Data Format.....	5-37
Table 5-50	MODE SELECT (6) Command Field Descriptions	5-37
Table 5-51	Initiator-Changeable Mode Pages.....	5-38
Table 5-52	Mode Page Types.....	5-39
Table 5-53	Mode Parameter List-Data Format.....	5-39
Table 5-54	Mode Parameter List-Field Descriptions	5-40
Table 5-55	Mode Parameter Header (6-Byte)-Data Format	5-40
Table 5-56	Mode Parameter Header- Field Descriptions.....	5-40
Table 5-57	Mode Parameter Block Descriptor-Data Format	5-40
Table 5-58	Mode Parameter Block Descriptor-Field Descriptions.....	5-41
Table 5-59	Categories of Changeable Pages.....	5-41
Table 5-60	Unit Attention Condition Page (Page 0).....	5-41
Table 5-61	Unit Attention Condition Page (Page 0).....	5-42
Table 5-62	Read-Write Error Recovery Page-(Page 1)	5-42
Table 5-63	Read-Write Error Recovery Page - Field Descriptions	5-43
Table 5-64	Disconnect-Reconnect Page (Page 2).....	5-44
Table 5-65	Disconnect-Reconnect-Field Description (Page 2)	5-44
Table 5-66	Verify Error Recovery Page-(Page 7)	5-45
Table 5-67	Verify Error Recovery Page-Field Descriptions (Page 7).....	5-46
Table 5-68	Caching Page-(Page 8)	5-46
Table 5-69	Caching Page-Field Descriptions	5-47

Table 5-70	Control Mode Page-Data Format (Page 0Ah)	5-48
Table 5-71	Control Mode Page-Field Descriptions.....	5-49
Table 5-72	Notch and Partition Page-(Page 0Ch).....	5-50
Table 5-73	Notch and Partition Page-Field Descriptions	5-51
Table 5-74	Port Control Page Short Format.....	5-52
Table 5-75	Port Control Page Long Format	5-53
Table 5-76	Margin Control Subpage (01h)	5-54
Table 5-77	Saved Training Configuration Subpage (02h).....	5-55
Table 5-78	Negotiated Settings Subpage (03h).....	5-56
Table 5-79	Transceiver Mode	5-56
Table 5-80	Report Transfer Capabilities Subpage (4).....	5-57
Table 5-81	Field Descriptions.....	5-57
Table 5-84	Information Exceptions Control Page-(Page 1Ch)	5-59
Table 5-85	Information Exceptions Control Page-Field Descriptions (Page 1Ch)	5-60
Table 5-86	Codes Used by the MRIE Field	5-61
Table 5-87	MODE SELECT (10) Command Descriptor Block-Data Format	5-62
Table 5-88	Mode Parameter Header (10-Byte)-Data Format	5-62
Table 5-89	Mode Parameter Block Descriptor-Data Format	5-62
Table 5-90	Mode Parameter Block Descriptor-Field Descriptions.....	5-63
Table 5-91	MODE SENSE (6) Command Descriptor Block-Data Format	5-64
Table 5-92	MODE SENSE Command-Field Descriptions	5-64
Table 5-93	Mode Parameter Header (6 Byte)-Data Format.....	5-64
Table 5-94	Mode Parameter Header and Block Descriptor-Field Descriptions	5-65
Table 5-95	Read-Only Mode Pages	5-65
Table 5-96	Format Device Page-(Page 3)	5-65
Table 5-97	Format Device Page-Field Descriptions (Page 3).....	5-66
Table 5-98	Rigid Disk Geometry Page-(Page 4).....	5-67
Table 5-99	Rigid Disk Geometry Page-Field Descriptions (Page 4)	5-68
Table 5-100	MODE SENSE (10) Command Descriptor Block-Data Format.....	5-69
Table 5-101	Mode Parameter Header (10 Byte)-Data Format.....	5-69
Table 5-102	Mode Parameter Block Descriptor-Data Format	5-69
Table 5-103	Mode Parameter Header & Block Descriptor-Field Descriptions.....	5-70
Table 5-104	PERSISTENT RESERVATION IN Command Descriptor Block-Data Format.....	5-71
Table 5-105	PERSISTENT RESERVATION IN Command-Field Descriptions	5-72
Table 5-106	Read Keys Parameters-Data Format.....	5-72
Table 5-107	Read Keys Parameters-Field Descriptions	5-73
Table 5-108	Read Reservations Parameters-Data Format	5-73
Table 5-109	Read Reservations Parameters-Field Descriptions	5-74
Table 5-110	PERSISTENT RESERVATION IN Read Reservations Descriptor-Data Format.....	5-74
Table 5-111	PERSISTENT RESERVATION IN Read Reservations Descriptor-Field Descriptions	5-75
Table 5-112	Persistent Reservation Type Codes and Their Meanings.....	5-76
Table 5-113	When Do Conflicts Between Existing Reservations and New Reservations Exist?	5-78

Table 5-114	PERSISTENT RESERVATION OUT Command Descriptor Block-Data Format.....	5-79
Table 5-115	PERSISTENT RESERVATION OUT Command-Field Descriptions.....	5-80
Table 5-116	PERSISTENT RESERVATION OUT Command's Service Action Descriptions	5-81
Table 5-117	Persistent Reservation Type Codes and Their Meanings.....	5-85
Table 5-118	PERSISTENT RESERVATION OUT Parameter List-Data Format.....	5-86
Table 5-119	PERSISTENT RESERVATION OUT Parameter List-Field Descriptions	5-87
Table 5-120	Device Server Interpretation of Service and Scope Value	5-88
Table 5-121	READ (6) Command-Data Format	5-89
Table 5-122	READ (6) Command-Field Descriptions.....	5-89
Table 5-123	READ (10) Command-Data Format	5-90
Table 5-124	READ (10) Command-Field Descriptions.....	5-90
Table 5-125	READ BUFFER Command-Data Format	5-91
Table 5-126	READ BUFFER Command-Field Descriptions	5-92
Table 5-127	READ CAPACITY Command-Data Format.....	5-93
Table 5-128	READ CAPACITY Command-Field Descriptions	5-93
Table 5-129	READ CAPACITY Returned Data-Data Format	5-93
Table 5-130	READ DEFECT DATA (10) Command-Data Format	5-94
Table 5-131	READ DEFECT DATA (10) Command-Field Description	5-95
Table 5-132	Defect Descriptor-Block Format.....	5-95
Table 5-133	Defect Descriptor-Bytes From Index Format, or Physical Sector Format....	5-96
Table 5-134	Defect List Header -Data Format.....	5-96
Table 5-135	Defect List Header-Field Descriptions.....	5-96
Table 5-136	READ DEFECT DATA (12) Command-Data Format	5-97
Table 5-137	READ DEFECT DATA (12) Command-Field Description	5-97
Table 5-138	Defect Descriptor-Block Format.....	5-98
Table 5-139	Defect Descriptor-Bytes From Index Format, or Physical Sector Format....	5-98
Table 5-140	Defect List Header-Data Format.....	5-98
Table 5-141	Defect List Header-Field Descriptions.....	5-98
Table 5-142	READ LONG Command Descriptor Block-Data Format	5-99
Table 5-143	READ LONG Command Descriptor Block-Field Descriptions.....	5-99
Table 5-144	READ LONG Command-Returned Data	5-99
Table 5-145	REASSIGN BLOCKS Command Descriptor Block-Data Format.....	5-100
Table 5-146	REASSIGN BLOCKS Defect List Header-Data Format.....	5-100
Table 5-147	REASSIGN BLOCKS Defect List Header-Field Description	5-100
Table 5-148	REASSIGN BLOCKS Defect Descriptor-Data Format	5-100
Table 5-149	Diagnostic Pages Supported by The Drives.....	5-101
Table 5-150	RECEIVE DIAGNOSTIC RESULTS Command Descriptor Block-Data Format.....	5-101
Table 5-151	RECEIVE DIAGNOSTIC RESULTS Command Descriptor Block-Field Descriptions	5-101
Table 5-152	Supported Diagnostics Pages Page-Data Format	5-102
Table 5-153	Translate Address Page-Data Format.....	5-103
Table 5-154	Translate Address Page-Field Descriptions	5-103
Table 5-155	RELEASE (6) Command Descriptor Block-Data Format	5-104

Table 5-156	RELEASE (6) Command-Field Descriptions	5-104
Table 5-157	RELEASE (10) Command Descriptor Block-Data Format	5-105
Table 5-158	RELEASE (10) Command-Field Descriptions	5-105
Table 5-159	REPORT DEVICE IDENTIFIER Command Descriptor Block-Data Format.....	5-106
Table 5-160	REPORT DEVICE IDENTIFIER Command Descriptor Block-Field Descriptions	5-106
Table 5-161	REPORT DEVICE IDENTIFIER Parameter List-Data Format.....	5-106
Table 5-162	REPORT DEVICE IDENTIFIER Parameter List-Field Descriptions.....	5-106
Table 5-163	REPORT LUNS Command Descriptor Block-Data Format.....	5-107
Table 5-164	REPORT LUNS Command Descriptor Block-Field Description	5-107
Table 5-165	LUN Reporting Parameter List -Data Format	5-107
Table 5-166	REQUEST SENSE Command Descriptor Block-Data Format	5-108
Table 5-167	REQUEST SENSE Command Descriptor Block-Field Description	5-108
Table 5-168	Sense Data Format for Error Code 70h or 71h-Data Format.....	5-109
Table 5-169	Sense Data Fields (Error Code 70h)-Field Descriptions	5-110
Table 5-170	Supported Sense Keys.....	5-111
Table 5-171	Sense Key Information Field Contents.....	5-112
Table 5-172	Supported Additional Sense Codes and Sense Code Qualifiers (in Hex) ...	5-113
Table 5-173	Sense-Key Specific Field Contents.....	5-121
Table 5-174	ILLEGAL REQUEST Sense Key Field Pointer Bytes\Data Format.....	5-122
Table 5-175	ILLEGAL REQUEST Sense Key Field Pointer Bytes\Field Descriptions	5-122
Table 5-176	NOT READY Sense Key - Progress Indication Bytes\Data Format	5-122
Table 5-177	NOT READY Sense Key - Progress Indication Bytes\Field Descriptions	5-122
Table 5-178	MEDIUM ERROR or RECOVERED ERROR Sense Key - Retry Count-Data Format.....	5-123
Table 5-179	MEDIUM ERROR or RECOVERED ERROR Sense Key - Retry Count	5-123
Table 5-180	RESERVE (6) Command Descriptor Block-Data Format	5-124
Table 5-181	RESERVE (6) Command-Field Descriptions	5-124
Table 5-182	RESERVE (10) Command Descriptor Block-Data Format	5-125
Table 5-183	RESERVE (10) Command-Field Descriptions	5-125
Table 5-184	Extent Descriptors-Data Format	5-126
Table 5-185	Reservation Types.....	5-126
Table 5-186	Parameter List When LongID and Extent Bits = 1-Data Format	5-126
Table 5-187	RESERVE (10) ID Only Parameter List-Data Format.....	5-126
Table 5-188	REZERO UNIT Command Descriptor Block-Data Format.....	5-127
Table 5-189	SEEK (6) Command Descriptor Block-Data Format.....	5-128
Table 5-190	SEEK (10) Command Descriptor Block-Data Format.....	5-129
Table 5-191	SEND DIAGNOSTIC Command Descriptor Block-Data Format	5-130
Table 5-192	SEND DIAGNOSTIC Command-Field Descriptions	5-130
Table 5-193	Supported Diagnostic Page List-Data Format	5-131
Table 5-194	Translate Address Page-Data Format.....	5-131
Table 5-195	SET DEVICE IDENTIFIER Command Descriptor Block-Data Format.	5-132
Table 5-196	SET DEVICE IDENTIFIER Command Descriptor Block-Field Descriptions	5-132

Table 5-197	SET DEVICE IDENTIFIER Parameter List-Data Format.....	5-132
Table 5-198	SET DEVICE IDENTIFIER Parameter List-Field Descriptions	5-132
Table 5-199	START STOP UNIT Command Descriptor Block-Data Format	5-133
Table 5-200	START STOP UNIT Command-Field Descriptions	5-133
Table 5-201	SYNCHRONIZE CACHE Command Descriptor Block-Data Format ..	5-134
Table 5-202	SYNCHRONIZE CACHE Command-Field Descriptions	5-134
Table 5-203	TEST UNIT READY Command Descriptor Block-Data Format	5-135
Table 5-204	VERIFY Command Descriptor Block-Data Format	5-136
Table 5-205	VERIFY Command-Field Descriptions.....	5-136
Table 5-206	WRITE (6) Command Descriptor Block-Data Format.....	5-137
Table 5-207	WRITE (6) Command-Field Descriptions.....	5-137
Table 5-208	WRITE (10) Command Descriptor Block-Data Format	5-138
Table 5-209	WRITE (10) Command-Field Descriptions.....	5-138
Table 5-210	WRITE AND VERIFY Command Descriptor Block-Data Format	5-139
Table 5-211	WRITE AND VERIFY Command -Field Descriptions.....	5-139
Table 5-212	WRITE BUFFER Command Descriptor Block-Data Format	5-140
Table 5-213	WRITE BUFFER Command -Field Descriptions.....	5-141
Table 5-214	WRITE LONG Command Descriptor Block-Data Format.....	5-142
Table 5-215	WRITE LONG Command-Field Descriptions	5-142
Table A-1	SCSI-2/SCSI-3 Equivalent Terminology.....	A-1
Table A-2	SCSI-3 Quick Reference – Commands.....	A-2
Table A-3	SCSI-3 Quick Reference – Messages.....	A-3
Table A-4	SCSI-3 Quick Reference - Pages	A-3
Table A-5	SCSI-3 Quick Reference – Sense Keys.....	A-5
Table A-6	SCSI-3 Quick Reference – Status Codes	A-5
Table B-1	Transfer Period Factor Field Values When PARL = 1	B-1
Table B-2	Transfer Period Factor Field Values When PARL = 0	B-2
Table B-3	Settings for the QAS, DT, and IU Fields of the Negotiated Rate Information Page (INQUIRY Command)	B-2

Chapter 1

ABOUT THIS MANUAL

This chapter gives an overview of the contents of this manual, including the intended audience, how the manual is organized, terminology and conventions, and references.

1.1 AUDIENCE

The Maxtor Atlas 10K V Product Manual is intended for reference by original equipment manufacturers (OEMs) that are integrating the disk drive into a system or subsystem, developers, and disk drive installers. Its primary audience is the OEM technical staff that makes disk drive purchase and configuration decisions, and system integrators that are responsible for the SCSI interface. This manual is not intended for end-users and is not a users manual or an installation guide. The manual provides information about installation, interface command implementation, maintenance, and gives the general specifications of the drive.

1.2 MANUAL ORGANIZATION

This manual is organized into the following chapters:

- Chapter 1 – About This Manual
- Chapter 2 – General Description
- Chapter 3 – Installation
- Chapter 4 – Specifications
- Chapter 5 – SCSI Description
- Chapter 6 – Feature Descriptions
- Appendix A – Quick Reference
- Appendix B – Negotiated Page Information Page Reference
- Glossary
- Index

1.3 TERMINOLOGY AND CONVENTIONS

In the Glossary at the back of this manual, you can find definitions for many of the terms used in this manual. In addition, the following abbreviations are used in this manual:

- ASIC application-specific integrated circuit
- Kbpi thousands of bits per inch
- dB decibels
- dBA decibels, A weighted
- ECC error correcting code
- Kfci thousands of flux changes per inch
- Gbit gigabit
- GB gigabyte
- Hz hertz
- KB kilobytes
- LSB least significant bit
- LVDS low voltage differential SCSI
- mA milliamperes
- MB megabytes (1 MB = 1,000,000 bytes when referring to disk transfer rates or storage capacities and 1,048,576 bytes in all other cases)
- Mbit/s megabits per second
- MB/s megabytes per second
- MHz megahertz
- ms milliseconds
- MSB most significant bit
- mV millivolts
- ns nanoseconds
- SCSI Small Computer System Interface
- tpi tracks per inch
- μ s microseconds
- V volts

The typographical and naming conventions used in this manual are listed below. Conventions that are unique to a specific table appear in the notes that follow that table.

Typographical Conventions:

- Names of Bits: Bit names are presented in initial capitals. An example is the Host Software Reset Bit.
- Commands: Firmware commands are listed as all capitals. An example is MODE SELECT.
- Parameters: Parameters are given as initial capitals when spelled out, and are given as all capitals when abbreviated. Examples are Prefetch Enable (PE) and Cache Enable (CE).
- Hexadecimal Notation: The hexadecimal notation is given in 9-point subscript form. An example is 30_H.
- Signal Negation: A signal that is non-active or is in its non-asserted state.
- Messages: A message that is sent from the drive to the host is listed in all capitals. An example is BUS DEVICE RESET.

Naming Conventions:

- Host: In general, the system in which the drive resides is referred to as the host.
- Computer Voice: This refers to items you type at the computer keyboard. These items are listed in 10-point, all capitals, Courier font. An example is FORMAT.

1.4 REFERENCES

For additional information about the SCSI interface, refer to:

- ANSI Small Computer System Interface-2 (SCSI-2) Specification, ANSI X3T9.2/86-109, Revision 10k.
- ANSI Small Computer System Interface-3 (SCSI-3) Specification, ANSI X3T10, August 1994.
- SPI-3 Specification Revision SPI-4 Spec Revision 10

Chapter 2

GENERAL DESCRIPTION

This chapter summarizes the general functions and key features of the Maxtor Atlas 10K V family of hard disk drives, as well as the applicable standards and regulations.

2.1 PRODUCT OVERVIEW

Maxtor Atlas 10K V hard disk drives are part of a family of high performance, 1-inch-high hard disk drives manufactured to meet the highest product quality standards. There are currently three models in the Maxtor Atlas 10K V series, with capacities of 73.5, 147.1, and 300 gigabytes (GB). These hard disk drives use nonremovable, 84 mm diameter hard disks and are available with the following SCSI configurations:

- Ultra320 SCSI, Ultra2, Ultra SCSI 68-pin Wide (16-bit)
- Ultra320 SCSI, Ultra2, Ultra SCSI 80-pin SCA-2 (16-bit)

The Maxtor Atlas 10K V hard disk drives feature an embedded SCSI drive controller and use SCSI commands to optimize system performance. Because the drive manages media defects and error recovery internally, these operations are fully transparent to the user.

The innovative design of the Maxtor Atlas 10K V hard disk drives enables Maxtor to produce a family of low-cost, high-reliability drives.

2.2 KEY FEATURES

The Maxtor Atlas 10K V hard disk drive includes the following key features:

General

- Formatted storage capacity of 73.5 GB (1 disk, 2 heads), 147.1 GB (2 disks, 4 heads), and 300 GB (4 disks, 8 heads)
- Low profile, 1-inch height
- Industry standard 3 1/2-inch form factor
- Embedded SCSI controller

Performance

- Average seek time of 4.0 ms (1D) , 4.2 (2D) and 4.4 (4D) for reads
- Average seek time of 4.5 ms (1D), 4.7 (2D), and 4.9 ms (4D) for writes
- 10,000 RPM rotational speed
- Average rotational latency of 3 ms
- 8 MB CACHE (ECC Protected) buffer. Look-ahead DisCache feature with continuous prefetch and WriteCache write-buffering capabilities
- Read-on-arrival firmware
- Tagged Command Queuing with Reordering
- ECC on-the-fly
- Highly automated SCSI protocol (including Auto Read/Write)
- 1:1 interleave on read/write operations
- High performance Ultra320 SCSI interface
- SCSI-2, and SCSI-3 supported
- Ultra 320 SCSI
- Ultra 160 SCSI
- Ultra-2 SCSI-3 LVDS transfer rates supported
- Fast Ultra and Ultra-2 SCSI transfer rates supported (SCSI-3 compliant)
- S.M.A.R.T. 2 (Self-Monitoring, Analysis and Reporting Technology)
- SCSI bus active negation drivers
- Burst data transfer rate of up to 6.0 MB/s asynchronous, 40 MB/s Ultra SCSI, 80 MB/s Ultra2 SCSI, 160 MB/s Ultra160 SCSI, and 320 MB/s Ultra320 SCSI.
- SCSI bus fairness

-
- MaXAdapt™
 - ~ Adaptive Active Filter (AAF)
 - ~ Rotational Vibration Compensation (RVC)
 - ~ Adaptive Bias Estimation (ABE)
 - ~ Virtual Cache Line (VCL)
 - Ultra 320 Performance Enhancements
 - ~ Double Transition (DT) Data Transfer
 - ~ Free Running Clock (FRC)
 - ~ Skew Compensation
 - ~ Cyclic Redundancy Check (CRC)
 - ~ Domain Validation
 - ~ Information Unit (IU) Transfers
 - ~ Transmitter Pre-Compensation with Cutback
 - ~ Quick Arbitration and Selection (QAS)
 - ~ Asynchronous Information Protection (AIP)
 - ~ SCSI Bus Fairness
 - ~ Flow Control

Reliability

- Automatic retry on read errors
- 360-bit, interleaved Reed-Solomon Error Correcting Code (ECC), with cross checking correction up to three separate bursts of 32 bits each totalling up to 96 bits in length
- Self-diagnostic firmware
- Transparent media defect mapping
- High performance, in-line defective sector skipping
- Reassignment of defective sectors discovered in the field, without reformatting
- Thermal Sensing monitors the drive temperature to ensure on-going drive reliability
- Shock Protection System III (SPS III) protects the drive against specific types of handling events that could cause damage to the drive.

Versatility

- Downloadable firmware
- Plug-and-Play SCSI
- SCSI-2, Ultra160 SCSI, and Ultra320 SCSI compatibility

2.3 REGULATORY COMPLIANCE STANDARDS

Maxtor Corporation's disk drive products meet all domestic and international product safety regulatory compliance requirements. Maxtor's disk drive products conform to the following specifically marked Product Safety Standards:

- Underwriters Laboratories (UL) Standard 1950 3rd Edition. This certificate is category UL recognized pertaining to all 3.5 inch series drives.
- Canadian Standards Association (CSA) Standard C.22.2 No. 950. This certificate is category c-UL recognized pertaining to all 3.5 inch series drives.
- TUV Rheinland Standard EN60950. This certificate is a category certification pertaining to all 3.5 inch series drives.

Product EMI/EMC Qualifications:

Maxtor Corporation's disk drive products meet all domestic and international electromagnetic emissions and immunity requirements. Maxtor's disk drive products conform to the following EMI/EMC Standards

- CE Mark (Europe) is a Self Declaration as per Directive 89/336, EN55022: 1998 (Emissions) and EN55024: 1998 (Immunity).
- C-Tick Mark (Australia/New Zealand) is a Self Declaration as per AS/NZS3548: 1998.
- BSMI Mark (Taiwan) is a Self Declaration as per CNS 13438:1998.
- MIC Mark (Korea) is a Certificate of registration for SCSI disk drives.
- Maxtor's disk drives are designed as a separate subassembly that conforms to the FCC Rules for Radiated and Conducted emissions, Part 15 Subpart J; Class B when installed in a given computer system.
- Maxtor drives conform to Canadian EMC Standard ICES-003.

2.4 HARDWARE REQUIREMENTS

The Maxtor Atlas 10K V hard disk drive is compatible with host computers and controllers that provide a 68-pin Wide, or 80-pin SCA-2 interface. A 50-pin to 68-pin adapter is required to use the 68-pin Wide drive in a 50-pin cabling configuration.

Termination is required on the Maxtor Atlas 10K V hard disk drives as they do not support on-board SCSI termination.

Chapter 3

INSTALLATION

This chapter explains how to unpack, configure, mount, and connect the Maxtor Atlas 10K V hard disk drive prior to operation. It also explains how to start up and operate the drive.

3.1 SAFETY, HANDLING, & ELECTROSTATIC DISCHARGE PROTECTION

3.1.1 Safety Precautions

For your safety, follow all safety procedures described here and in other sections of the manual.

- Remove power from the computer system (or expansion unit) before installing or removing the drive to prevent the possibility of electrical shock or damage to the drive. Unplug the unit containing the drive to provide an added measure of safety.
- Read, understand, and observe all label warnings.

3.1.2 Handling

Damage to the drive can occur as the result of careless handling, vibration, shock, or electrostatic discharge (ESD). Always handle the drive with care to avoid damage to the precision internal components.

CAUTION: A 1/4-inch drop onto a hard surface can damage the drive.

Follow these guidelines to avoid damage to the drive:

- Always observe prescribed ESD precautions.
- Keep the drive in its anti-static bag until ready to install.
- Always use a properly fitted wrist strap or other suitable ESD protection when handling the drive.
- Hold drive only by its sides. Do not touch any components on the PCBA.

-
- Always handle the drive carefully and gently. A drop of 1/4 inch onto a bench or desktop can damage a drive.
 - Do not bump, jar, or drop the drive. Use care when transporting the drive.
 - Always gently place the drive flat, PCB side down, on an appropriate ESD-protected work surface to avoid the drive being accidentally knocked over.
 - Do not pack other materials with the drive in its shielded bag.
 - Place the drive in the anti-static bag before placing in shipping container.
 - Do not stack objects on the drive.
 - Never force the drive or the mounting brackets into the drive bay.
 - Do not expose the drive to moisture.
 - Do not damage any seals on the drive; doing so may void the warranty.

3.1.3 Electrostatic Discharge (ESD) Protection

Various electrical components within the disk drive are sensitive to static electricity and Electrostatic Discharge (ESD). Even a static buildup or discharge that is too slight to feel can be sufficient to destroy or degrade a component's operation.

To minimize the possibility of ESD-related damage to the drive, we strongly recommend using both, a properly installed workstation anti-static mat and a properly installed ESD wrist strap. When correctly installed, these devices reduce the buildup of static electricity which might harm the drive.


- Observe the following precautions to avoid ESD-related problems:
- Use a properly installed anti-static pad on your work surface.
- Always use a properly fitted and grounded wrist strap or other suitable ESD protection when handling the drive and observe proper ESD grounding techniques.
- Hold the drive only by its sides. Do not touch any components on the PCBA.
- Leave the drive in its anti-static bag until you are ready to install it in the system.
- Place the drive on a properly grounded anti-static work surface pad when it is out of its protective anti-static bag.
- Do not use the bag as a substitute for the work surface anti-static pad. The outside of the bag may not have the same anti-static properties as the inside. It could actually increase the possibility of ESD problems.
- Do not use any test equipment to check components on the electronics module. There are no user-serviceable components on the drive.

3.2 SPACE REQUIREMENTS

The Maxtor Atlas 10K V hard disk drive is shipped without a faceplate and comes in the following SCSI interface configurations:

- 68-pin Wide SCSI
- 80-pin SCA-2 SCSI

Figure 3-1 Shows The Mechanical Dimensions of the drives.



3.2.1 Shock Feet

Maxtor Atlas 10K V hard disk drives are outfitted with plastic shock feet on the bottom edge of the base casting, near the corners, beneath the side mounting holes (translucent), and near the corners of the top cover next to the screws (black). The shock feet give an additional level of isolation to prevent the head and disk damage that occasionally occurs during unpacking, staging, and installation. The shock feet attenuate the short-pulse shocks that occur when placing the drive on a hard surface. If the drive is tested on a hard surface, it should be supported such that the shock feet are not in contact with a hard surface (the drive should be supported in the middle, between the shock feet).

Note: To provide optimal protection the shock feet are designed to exceed the form factor when uncompressed.

3.3 UNPACKING INSTRUCTIONS

CAUTION: The maximum limits for physical shock can be exceeded if the drive is not handled properly. Special care should be taken not to bump or drop the drive.

1. Open the shipping container and remove the packing assembly that contains the drive.
2. Remove the drive from the packing assembly.

CAUTION: During shipment and handling, the antistatic electrostatic discharge (ESD) bag prevents electronic component damage due to electrostatic discharge. To avoid accidental damage to the drive, do not use a sharp instrument to open the ESD bag. Save the packing materials for possible future use.

3. When you are ready to install the drive, remove it from the ESD bag.

Figure 3-2 shows the packing assembly for a single Maxtor Atlas 10K V hard disk drive. Shipping containers for 25-pack are available for multiple drive shipments.

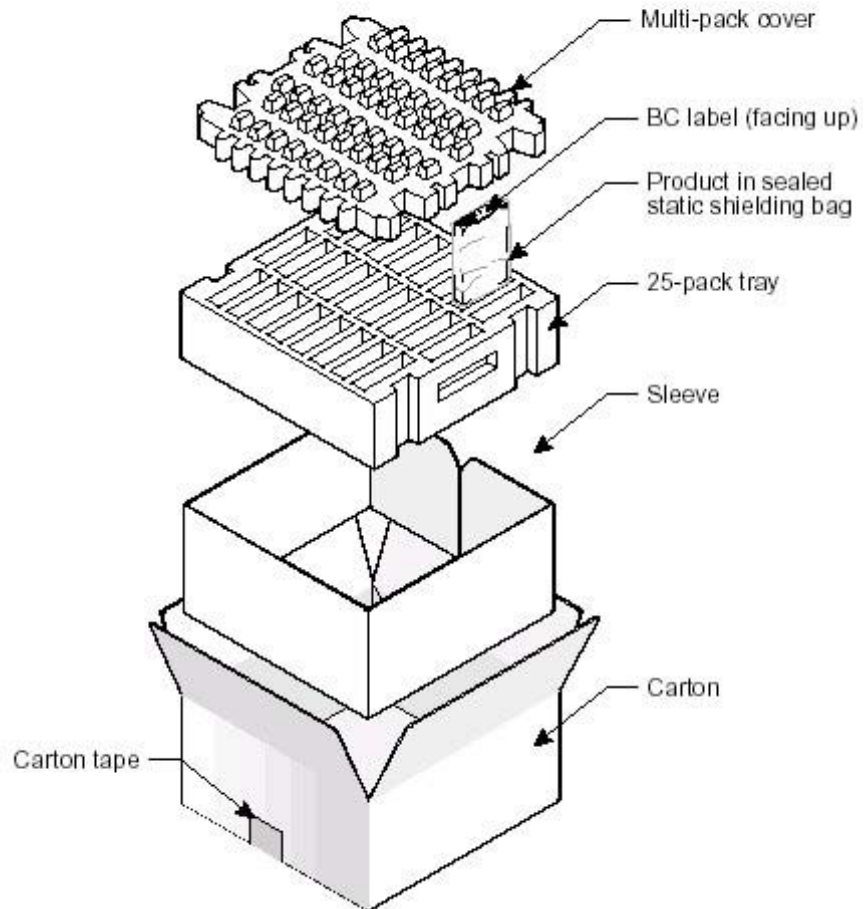


Figure 3-2 Drive Packing Assembly

3.4 HARDWARE OPTIONS

3.4.1 Configuration Jumpers and Connections

This section includes setup and configuration information for Maxtor Atlas 10K V drives. These disk drives include

- The 16-bit multimode Ultra320 SCSI, wide version with 68-pin SCSI connector,
- The 16-bit multimode Ultra320 SCSI, version with SCA-2 80-pin connector.

Specific individual settings for each drive type are described in Sections 3.5.1 through 3.5.4.

CAUTION: Before you begin, review the Safety, ESD, and Handling precautions described at the beginning of this manual to avoid personal injury or damage to equipment.

3.4.2 Jumper Options on the 68-Pin Wide PCB

This section describes how to configure the jumpers on Maxtor Atlas 10K V disks with 68-pin SCSI interface connectors. The following features are jumper-selectable:

- SCSI ID (0), (1), (2), (3) – SCSI Bus Device Identification
- Delay Spin
- Single-Ended, Disable LVD
- Busy Out (Remote LED)
- Write Protect
- Stagger Spin

Note: The disk drive does not support on-board SCSI termination.

Note: The configuration of a Maxtor Atlas 10K V hard disk drive depends on the host system in which it is to be installed. Figure 3-3 shows the printed circuit board (PCB) assemblies for 68-pin SCSI configurations, indicating the jumpers that control some of these options.




Figure 3-3 Jumper Locations on the 68-Pin Wide SCSI Drive PCB

Table 3-1 SCSI ID Selection on Option Connector (68-Pin SCSI Connector Drives)

SCSI ID	Jumper Location – J3 Option Header			
	Pin Pair 7/8	Pin Pair 5/6	Pin Pair 3/4	Pin Pair 1/2
0	0	0	0	0
1	1	0	0	0
2	0	1	0	0
3	1	1	0	0
4	0	0	1	0
5	1	0	1	0
6	0	1	1	0
7	1	1	1	0
8	0	0	0	1
9	1	0	0	1
10	0	1	0	1
11	1	1	0	1
12	0	0	1	1
13	1	0	1	1
14	0	1	1	1
15	1	1	1	1

Note: 0 = No Jumper, 1 = Jumper Installed

Configure the drive for remote (external) SCSI ID selection by removing the SCSI ID jumpers (if present) from the referenced SCSI ID pins. Then connect the leads from the external selection switch to the referenced pins. Observe the following guidelines while doing so:

- ID bit 0, at Pin 8, is the Least Significant Bit.
- SCSI ID bits 0, 1, 2, and 3 (pins 8, 6, 4, and 2, respectively) are active LOW signals. That is, the bit is a 1 if the corresponding remote switch is closed to ground or jumper installed.
- Use pins 1, 3, 5 and 7 as the associated ground returns for ID bits 3, 2, 1, and 0, respectively.

3.4.2.1 Write Protection

To configure Write Protection for the drive, install a jumper across pin pair 11/12 on the J3 Option Header. To disable Write Protection on the drive, remove the jumper.

3.4.2.2 Delay Spin (DS), Stagger Spin (SS)

Maxtor Atlas 10K V drives have three Spin Up modes:

Option 1 (No jumpers installed):

Spin up immediately when power is applied. Verify that no jumper is installed across the Delay Spin pin pair of the J3 Option Header.

Option 2 (Delay Spin jumper installed):

Spin up on START STOP UNIT command: Install the jumper across pin pair 15/16 (GND/Delay Spin) on the J3 Option Header.

3.4.2.3 Single-Ended Operation – Force SE (Disable LVD)

Install a jumper across pin pair 17/18 (SE) on the J3 Option Header to operate the disk drive as a single-ended device. Remove the SE jumper for LVD operation and monitoring of the DIFFSENS signal.

3.4.2.4 Remote Busy and Fault Displays

Busy and Fault status of the drive can be monitored remotely by connecting a remote (external) Busy and/or remote Fault display LEDs.

Remote Fault LED

On the J1 Option Connector, connect the cathode side of the remote Fault LED to pin 2, Fault LED. Connect the anode side of the LED to pin 11, +5V.

3.4.3 SCA-2 80-Pin Connector Versions

This section describes the SCA-2 (Single Connector Attachment) 80-pin connector for Maxtor Atlas 10K V drives with the following features:

- SCSI ID
- Spin Up
- Activity LED displays

Use Figure 3-4 to locate the appropriate pins for configuring the drive. Note that Figure 3-4 does not call out each of the 80 pins on the connector, but rather illustrates the layout of the pins.

Note: The SCA-2 Connector version of the disk drives does not provide the following jumper configuration: TERMPWR, Active Termination, or Write Protection.

3.4.3.1 Termination for the 80-Pin SCA-2 Connector Versions

CAUTION: These versions of Maxtor Atlas 10K V disk drives cannot be configured to provide bus termination. Therefore, be sure to properly terminate the SCSI bus on which this drive is installed.

Note: Refer to your system or SCSI controller documentation regarding any additional recommendations regarding drive placement on the SCSI bus and SCSI bus termination.

3.4.3.2 SCSI ID for SCA-2 Versions

Each SCSI device on the bus must have a unique SCSI ID number assigned to it. The drive can be configured for SCSI ID numbers that range from 0 through 15.

Configure the SCSI ID by providing the proper open or ground signal inputs to the referenced pins of the drive's 80-pin new version SCA-2 connector (Figure 3-4). Refer to Table 3-2 for SCSI ID pin assignments.

Note: Refer to your system or SCSI controller documentation for specific recommendations about assigning SCSI ID numbers for your specific system.




Figure 3-4 Pin Locations on SCA-2 Connector

Table 3-2 SCSI ID Pin Assignments (SCA-2 Connector Versions of the Disk Drive)

SCSI ID	Location on SCA Connector			
	ID3 - Pin 80	ID2 - Pin 40	ID1 - Pin 79	ID0 - Pin 39
0	0	0	0	0
1	0	0	0	1
2	0	0	1	0
3	0	0	1	1
4	0	1	0	0
5	0	1	0	1
6	0	1	1	0
7	0	1	1	1
8	1	0	0	0
9	1	0	0	1
10	1	0	1	0
11	1	0	1	1
12	1	1	0	0
13	1	1	0	1
14	1	1	1	0
15	1	1	1	1

Note: 0 indicates an open circuit and 1 indicates ground.

3.4.3.3 Spin Up for the 80-Pin SCA-2 Connector Versions

Maxtor Atlas 10K V drives have three Spin Up modes:

- Option 1** Spin up occurs immediately when power is applied.
- Option 2** Drive spin up occurs after a predetermined delay following power on. Set the delay parameters with the MODE SELECT Command, Maxtor (Vendor) Special Function Control Page (39h). The delay is equal to a user-specified multiplier multiplied by the numerical SCSI ID of the drive. This will give a staggered spin-up in multiple-drive installations.
- Option 3** Drive spin up is controlled by the START STOP UNIT command.

Configure the desired spin up option (Table 3-3) by setting the state of the DELAY_SPIN (Pin 38) and STAGGER_SPIN (Pin 78) inputs on the 80-pin SCA-2 connector (Figure 3-4). The states of these signals are set by using either hard-wired connections at the backplane or backplane logic.

Table 3-3 Spin Up on Power On Options

Option	STAGGER_SPIN (Pin 78)	DELAY_SPIN (Pin 38)
<i>Option 1</i> – Spin Up When Power is Applied	Open	Open
<i>Option 2</i> – Spin Up After Delay	Ground	Open
<i>Option 3</i> – Spin Up on START Command	Open	Ground
Reserved	Ground	Ground

3.4.3.4 Activity LED for SCA-2 Connector Versions

The drive provides the output BUSY_OUT signal to power a user-supplied activity LED.


The output indicates the drive is performing a SCSI operation. To use this output, connect a user-supplied LED cathode to the BUSY_OUT connection, pin 77 on the SCA Connector (Figure 3-4). The LED anode must be attached to the proper +5 VDC supply through an appropriate current-limiting resistor.


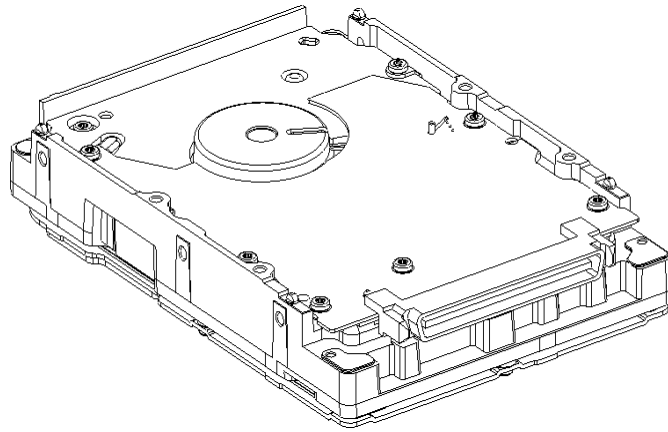
3.5 INTERFACE CONNECTOR (J1)

The configuration of J1 is different for the 68-pin and 80-pin SCSI variations.

Figure 3-5 shows the various connector styles.

Figure 3-5 J1 Interface Connector Configurations





3.5.1 68-Pin Wide SCSI Connector - LVD

Table 3-4 68-Pin Wide LVD Pin Assignments

SIGNAL NAME	CONNECTOR CONTACT NUMBER	CABLE CONDUCTOR NUMBER		CONNECTOR CONTACT NUMBER	SIGNAL NAME
+DB (12)	1	1	2	35	-DB (12)
+DB (13)	2	3	4	36	-DB (13)
+DB (14)	3	5	6	37	-DB (14)
+DB (15)	4	7	8	38	-DB (15)
+DB (P1)	5	9	10	39	-DB (P1)
+DB (0)	6	11	12	40	-DB (0)
+DB (1)	7	13	14	41	-DB (1)
+DB (2)	8	15	16	42	-DB (2)
+DB (3)	9	17	18	43	-DB (3)
+DB (4)	10	19	20	44	-DB (4)
+DB (5)	11	21	22	45	-DB (5)
+DB (6)	12	23	24	46	-DB (6)
+DB (7)	13	25	26	47	-DB (7)
+DB (P)	14	27	28	48	-DB (P)
GROUND	15	29	30	49	GROUND
DIFFSENS	16	31	32	50	GROUND
TERMPWR	17	33	34	51	TERMPWR
TERMPWR	18	35	36	52	TERMPWR
RESERVED	19	37	38	53	RESERVED
GROUND	20	39	40	54	GROUND
+ATN	21	41	42	55	-ATN
GROUND	22	43	44	56	GROUND
+BSY	23	45	46	57	-BSY
+ACK	24	47	48	58	-ACK
+RST	25	49	50	59	-RST
+MSG	26	51	52	60	-MSG
+SEL	27	53	54	61	-SEL
+C/D	28	55	56	62	-C/D
+REQ	29	57	58	63	-REQ
+I/O	30	59	60	64	-I/O
+DB (8)	31	61	62	65	-DB (8)
+DB (9)	32	63	64	66	-DB (9)
+DB (10)	33	65	66	67	-DB (10)
+DB (11)	34	67	68	68	-DB (11)

Note: The conductor number refers to the conductor position when using 0.635 mm (0.025 inch) centerline flat ribbon cable. Other cable types may be used to implement equivalent contact assignments.

3.5.1.1 68-Pin Wide Mating Connector

The cable plug connector that mates with the 68-pin hard disk drive connector is the AMP AMPLIMITE .050 Series III, part number 749925-5.

3.5.2 80-Pin SCA-2 SCSI Connector - LVD

Table 3-5 80-Pin SCA-2 LVD Pin Assignments

80-PIN CONNECTOR CONTACT AND SIGNAL NAME		CABLE CONDUCTOR NUMBER	80-PIN CONNECTOR CONTACT AND SIGNAL NAME	
12V CHARGE	1	NOT APPLICABLE	12V GROUND	41
12V	2		12V GROUND	42
12V	3		12V GROUND	43
12 V	4		MATED 1	44
OPT 3.3 VOLTS	5		OPT 3.3V CHARGE	45
OPT 3.3 VOLTS	6		DIFFSNS	46
- DB (11)	7		+DB (11)	47
- DB (10)	8		+DB (10)	48
- DB (9)	9		+DB (9)	49
- DB (8)	10		+DB (8)	50
- I/O	11		+I/O	51
- REQ	12		+REQ	52
- C/D	13		+C/D	53
- SEL	14		+SEL	54
- MSG	15		+MSG	55
- RST	16		+RST	56
- ACK	17		+ACK	57
- BSY	18		+BSY	58
- ATN	19		+ATN	59
- DB (P0)	20		+DB (P0)	60
- DB (7)	21		+DB (7)	61
- DB (6)	22		+DB (6)	62
- DB (5)	23		+DB (5)	63
- DB (4)	24		+DB (4)	64
-DB (3)	25		+DB (3)	65

-DB (2)	26	NOT APPLICABLE	+DB (2)	66
-DB (1)	27		+DB (1)	67
-DB (0)	28		+DB (0)	68
-DB (P1)	29		+DB (P1)	69
-DB (15)	30		+DB (15)	70
-DB (14)	31		+DB (14)	71
-DB (13)	32		+DB (13)	72
-DB (12)	33		+DB (12)	73
5V	34		MATED 2	74
5V	35		5V GROUND	75
5V CHARGE	36		5V GROUND	76
SPINDLE SYNC	37		ACTIVE LED OUT	77
RMT START	38		DLYD_START	78
SCSI ID (0)	39		SCSI ID (1)	79
SCSI ID (2)	40	SCSI ID (3)	80	

3.5.2.1 80-Pin SCA-2 Mating Connectors

The compatible 80-pin mating connectors are:

- Right-angle receptacle, AMP CHAMP, .050 Series I, part number 787535-1.
- Vertical receptacle, AMP CHAMP .050 Series I, part number 787311-2.
- Vertical receptacle, BERG part number 71780-001.

3.5.3 68-Pin Wide Single-Ended SCSI Connector

Table 3-6 68-Pin Wide Single-Ended Pin Assignments

SIGNAL NAME	CONNECTOR CONTACT NUMBER	CABLE CONDUCTOR NUMBER		CONNECTOR CONTACT NUMBER	SIGNAL NAME
GROUND	1	1	2	35	-DB (12)
GROUND	2	3	4	36	-DB (13)
GROUND	3	5	6	37	-DB (14)
GROUND	4	7	8	38	-DB (15)
GROUND	5	9	10	39	-DB (P1)
GROUND	6	11	12	40	-DB (0)
GROUND	7	13	14	41	-DB (1)
GROUND	8	15	16	42	-DB (2)
GROUND	9	17	18	43	-DB (3)
GROUND	10	19	20	44	-DB (4)
GROUND	11	21	22	45	-DB (5)
GROUND	12	23	24	46	-DB (6)
GROUND	13	25	26	47	-DB (7)
GROUND	14	27	28	48	-DB (P)
GROUND	15	29	30	49	GROUND
GROUND	16	31	32	50	GROUND
RESERVED	19	37	38	53	RESERVED
GROUND	20	39	40	54	GROUND
GROUND	21	41	42	55	-ATN
GROUND	22	43	44	56	GROUND
GROUND	23	45	46	57	-BSY
GROUND	24	47	48	58	-ACK
GROUND	25	49	50	59	-RST
GROUND	26	51	52	60	-MSG
GROUND	27	53	54	61	-SEL
GROUND	28	55	56	62	-C/D
GROUND	29	57	58	63	-REQ
GROUND	30	59	60	64	-I/O
GROUND	31	61	62	65	-DB (8)
GROUND	32	63	64	66	-DB (9)
GROUND	33	65	66	67	-DB (10)
GROUND	34	67	68	68	-DB (11)

Note:

1. The minus sign next to a signal indicates active low.
2. The conductor number refers to the conductor position when using 0.635 mm (0.025 inch) centerline flat ribbon cable. Other cable types may be used to implement equivalent contact assignments.

3.5.4 80-Pin SCA-2 Single-Ended SCSI Connector

Table 3-7 80-Pin SCA-2 Single-Ended Pin Assignments

80-PIN CONNECTOR CONTACT AND SIGNAL NAME		CABLE CONDUCTOR NUMBER	80-PIN CONNECTOR CONTACT AND SIGNAL NAME	
12V CHARGE	1	NOT APPLICABLE	12V GROUND	41
12V	2		12V GROUND	42
12V	3		12V GROUND	43
12 V	4		MATED 1	44
OPT 3.3 VOLTS	5		OPT 3.3V CHARGE	45
OPT 3.3 VOLTS	6		GROUND	46
-DB (11)	7		GROUND	47
-DB (10)	8		GROUND	48
-DB (9)	9		GROUND	49
-DB (8)	10		GROUND	50
-I/O	11		GROUND	51
-REQ	12		GROUND	52
-C/D	13		GROUND	53
-SEL	14		GROUND	54
-MSG	15		GROUND	55
-RST	16		GROUND	56
-ACK	17		GROUND	57
-BSY	18		GROUND	58
-ATN	19		GROUND	59
-DB (P0)	20		GROUND	60
-DB (7)	21		GROUND	61
-DB (6)	22		GROUND	62
-DB (5)	23		GROUND	63
-DB (4)	24		GROUND	64
-DB (3)	25		GROUND	65
-DB (2)	26		GROUND	66
-DB (1)	27		GROUND	67
-DB (0)	28		GROUND	68
-DB (P1)	29		GROUND	69
-DB (15)	30		GROUND	70
-DB (14)	31		GROUND	71
-DB (13)	32		GROUND	72
-DB (12)	33		GROUND	1
5V	34		MATED 2	74
5V	35		5V GROUND	75
5V CHARGE	36		5V GROUND	76
SPINDLE SYNC	37		ACTIVE LED OUT	77
RMT START	38		DLYD_START	78
SCSI ID (0)	39		SCSI ID (1)	79
SCSI ID (2)	40		SCSI ID (3)	80

3.6 DRIVE MOUNTING AND INSTALLATION

Drive mounting orientation, clearance, and ventilation requirements are described in the following subsections.

3.6.1 Orientation

The mounting holes on the Maxtor Atlas 10K V hard disk drive allow the drive to be mounted in any orientation. Figure 3-6 shows the location of the three mounting holes on each side of the drive. The drives also can be mounted using the four mounting hole locations on the PCB side.

All dimensions are in millimeters. For mounting, #6-32 UNC screws are recommended.




Figure 3-6 Mounting Dimensions Maxtor Atlas 10K V Drives

3.6.2 Mounting Screw Clearance

The printed-circuit board assembly (PCBA) is very close to the mounting holes. Figure 3-7 specifies the clearance between the screws in the mounting holes and the PCBA. Do not use mounting screws longer than the maximum lengths specified in Figure 3-7. The specified screw length allows full use of the mounting-hole threads, while avoiding damaging or placing unwanted stress on the PCBA. Clearance from the drive to any other surface (except mounting surfaces) must be a minimum of 1.25 mm (0.05 inches).




Figure 3-7 Mounting Screw Clearance for Maxtor Atlas 10K V Drive

CAUTION: The PCB is very close to the mounting holes. Do not exceed the specified length for the mounting screws. The specified screw length allows full use of the mounting-hole threads, while avoiding damaging or placing unwanted stress on the PCB. Figure 3-7 specifies the minimum clearance between the PCB and the screws in the mounting holes. The maximum torque applied to the screws must not exceed 8.6 inch-pounds.

3.6.3 Mounting

For the best results during performance benchmark testing, it is highly recommended that the drive be mounted firmly in a system or fixture, rather than sitting unconstrained on a tabletop. If it is necessary to do early testing of the drive while it is unconstrained, the drive should rest on a flat, smooth, semi-cushioned surface (similar to a mousepad). The drive should not be operated on a hard surface—this avoids performance degradations due to the higher incidence of recovered errors.

3.6.4 Ventilation

The Maxtor Atlas 10K V hard disk drive may operate with or without a cooling fan provided that the maximum HDA temperature as measured on the drive's top cover does not exceed 131°F (60°C). Drive reliability and warranty will be limited if the drive is exposed to temperatures greater than 60°C. Figures 3-8 and 3-9 show airflow recommended for adequate cooling. Clearance from the drive to any other surface above and below the drive must be a minimum of 1.25mm (1/16 inch). Maxtor leaves the design and application of cooling and clearance for the disk drive to the User, but the drive must maintain a case temperature at or below 60°C.




Figure 3-8 Lengthwise Airflow Cooling




Figure 3-9 Crosswise Airflow Cooling

Chapter 4 SPECIFICATIONS

This chapter gives a detailed description of the physical, electrical, and environmental characteristics of the Maxtor Atlas 10K V hard disk drives.

4.1 SPECIFICATION SUMMARY

Table 4-1 Specifications

DESCRIPTION	MAXTOR ATLAS 10K V 73.5 GB	MAXTOR ATLAS 10K V 147.1 GB	MAXTOR ATLAS 10K V 300 GB
Non-RoHS Model Numbers:	8D073J0 8D073L0	8D147J0 8D147L0	8D300J0 8D300L0
RoHS Model Numbers:	8J073J0 8J073L0	8J147J0 8J147L0	8J300J0 8J300L0
Formatted Capacity	73.5 GB	147.1 GB	300 GB
Nominal rotational speed (rpm)	10,000	10,000	10,000
Number of Disks	1	2	4
Number of R/W heads	2	4	8
Data Organization:			
Maximum BPI	732K	732K	732K
Zones per surface	16	16	16
Track density	102,000 TPI	102,000 TPI	102,000 TPI
User Data Cylinders	81,782	81,782	81,782
Sectors per track:			
Inside zone	624	624	624
Outside zone	1114	1114	1114
Total User Sectors	143,666,191	287,332,383	585,937,500
Bytes per sector	512 - 524	512 - 524	512 - 524
Recording technology	PRML	PRML	PRML
Maximum linear density	678 fci	678 fci	678 fci

DESCRIPTION	MAXTOR ATLAS 10K V 73.5 GB	MAXTOR ATLAS 10K V 147.1 GB	MAXTOR ATLAS 10K V 300 GB
Maximum effective areal density	75 Gbits/in ²	75 Gbits/in ²	75 Gbits/in ²
Servo Mechanical Timing Specifications:			
Sequential Head Switch Time ²	0.3 ms read 0.5 ms write	0.3 ms read 0.5 ms write	0.3 ms read 0.5 ms write
Sequential Cylinder Switch Time ³	0.3 ms read 0.5 ms write	0.3 ms read 0.5 ms write	0.3 ms read 0.5 ms write
Random ⁴ Average Seek (Read)	4.0 ms typical	4.2 ms typical	4.4 ms typical
Random ⁴ Average Seek (Write)	4.5 ms typical	4.7 ms typical	4.9 ms typical
Full Stroke ⁵ Seek (Read)	11.0 ms typical	11.0 ms typical	11.0 ms typical
SCSI 'Hard' Reset Time ⁶	22 ms typical 25 sec maximum	22 ms typical 25 sec maximum	22 ms typical 25 sec maximum
Power on to Drive Ready ⁷	25 sec typical	25 sec typical	25 sec typical
Data transfer Rates:			
Sequential Throughput Read/Write	89 MB/s maximum	89 MB/s maximum	89 MB/s maximum
Buffer Size	8 MB	8 MB	8 MB
Reliability:			
Seek error rate ⁸	<10 in 10 ⁸	<10 in 10 ⁸	<10 in 10 ⁸
Unrecoverable error rate ⁸	<1 in 10 ¹⁵	<1 in 10 ¹⁵	<1 in 10 ¹⁵
Error correction method (with cross check)	48 Bytes Reed-Solomon with XCHK	48 Bytes Reed-Solomon with XCHK	48 Bytes Reed-Solomon with XCHK
Contact Start/Stop Cycles	50,000 min.	50,000 min.	50,000 min.
Auto head-park method	Magnetic Latch	Magnetic Latch	Magnetic Latch

Note: Seek time is defined as the time required for the actuator to seek to a new position and settle on a track. It is measured by averaging the execution time of a minimum of 1000 operations of the indicated type as shown in Table 4-1. The seek times include head settling time, but do not include command overhead time, time spent reading or writing data, or rotational latency delays. Unless otherwise specified, read/seek command settling criteria is assumed.

1. Typical specifications assume 25°C ambient temperature, nominal supply voltages and no applied shock or vibration.
Maximum specifications assume worst case extremes of operating temperature, humidity, and supply voltages.
2. Sequential Head Switch time is the time from the conclusion of the last sector of a track to the beginning of the first logical sector on the next track of the same cylinder. It includes sequencer overhead for write setup on head and cylinder switch.
3. Sequential Cylinder Switch time is the time from the conclusion of the last sector of a cylinder to the first logical sector on the next cylinder. It includes sequencer overhead for write setup on head and cylinder switch.
4. Random LBA.
5. Full Stroke seek is defined as a seek from cylinder 0 to maximum cylinder or vice versa. It may include one head switch.
6. SCSI 'Hard' Reset time is the time from Reset to Selection.
7. At power on start-up error algorithms are used. These recovery routines may extend the time to Drive Ready by as much as 30 seconds.
8. Refer to Section 4.11, "DISK ERRORS" for details on error rate definitions.

4.2 FORMATTED CAPACITY

At the factory, the Maxtor Atlas 10K V receives a low-level format that creates the actual tracks and sectors on the drive. Table 4-2 shows the capacity resulting from this process. Formatting done at the user level for operation with DOS, UNIX, or other operating systems, will result in less capacity than the physical capacity shown.

Table 4-2 Formatted Capacity

	MAXTOR ATLAS 10K V 73.5 GB	MAXTOR ATLAS 10K V 147.1 GB	MAXTOR ATLAS 10K V 300 GB
Formatted Capacity	73.5 GB	147.1 GB	300 GB
Number of 514-524 byte sectors available	139,463,602	279,041,740	566,793,894
512 byte Sectors	143,666,191	287,332,383	585,937,500

4.3 DATA TRANSFER RATES

Data is transferred from the disk to the read buffer at a rate of up to 89 MB/s in bursts. Data is transferred from the read buffer to the SCSI bus at a rate of up to 6 MB/s in the asynchronous mode, or at up to 320 MB/s in the synchronous mode. For more detailed information on interface timing, refer to Chapter 6.

4.4 TIMING SPECIFICATIONS

Table 4-3 illustrates the timing specifications of the Maxtor Atlas 10K V hard disk drive.

Table 4-3 Timing Specifications for Maxtor Atlas 10K V

PARAMETER	MAXTOR ATLAS 10K V 73.5/147.1/300 GB	
	TYPICAL NOMINAL ¹	MAXIMUM ¹
Sequential Head Switch Time ²	0.3 ms typical (read) 0.5 ms typical (write)	N/A
Sequential Cylinder Switch Time ³	0.3 ms typical (read) 0.5 ms typical (write)	N/A
Random Average Seek (Read) ⁴	<4.0/ 4.2 / 4.4 ms	4.8 ms
Random Average Seek (Write) ⁴	4.5 / 4.7/ 4.9 ms	5.3 ms
Average Rotational Latency	3 ms	3 ms
1/3 Stroke Seek (Read) ⁵	<3 ms	4.5 ms
Full-Stroke Seek ⁶	11.0 ms	12.0 ms
SCSI "Hard Reset Time" ⁷	20 ms	20 ms
Power On to Drive Ready ⁸	25 seconds	30 seconds
Power On to Selection	2.5 seconds	3 seconds

1. Typical specifications assume 25°C ambient temperature, nominal supply voltages, and no applied shock or vibration. Maximum specifications assume worst case extremes of operating temperature, humidity, and supply voltages.
2. Sequential Head Switch Time is the time from the conclusion of the last sector of a track to the beginning of the first logical sector on the next track of the same cylinder.
3. Sequential Cylinder Switch Time is the time from the conclusion of the last sector of a cylinder to the first logical sector on the next cylinder.
4. Random LBA.
5. 1/3 Stroke Seek is defined as any seek of maximum length cylinder/3. 1/3 stroke seek may include one head switch.
6. Full Stroke Seek is defined as a seek from cylinder 0 to maximum cylinder or vice versa. Full stroke may include one head switch.
7. SCSI "Hard Reset Time" is the time from Reset to Selection.

-
8. At power on start up error algorithms are used and may extend the time to Drive Ready to as long as 30 seconds.

4.5 POWER

The Maxtor Atlas 10K V hard disk drive operates from two supply voltages:

MAXTOR ATLAS 10K IV 73.5/147.1/300 GB	
+12 V	+/-10% @ Spin up, +/- 5% While running
+5 V	+/- 5%

Allowable ripple and noise for each voltage:

+12 V	800 mV p-p (100 Hz to 8 KHz) 450 mV p-p (8 KHz to 20 KHz) 250 mVp-p (20KHz - 20MHz)
+5 V	250 mV p-p (100 Hz to 20MHz)

4.5.1 Power Sequencing

You may apply the power in any order or manner, or open either the power or power return line with no loss of data or damage to the disk drive. However, data may be lost in the sector being written at the time of power loss. The drive can withstand transient voltages of +10% to -10% from nominal while powering up or down.

4.5.2 Power Reset Limits

When powering up, the drive remains reset (inactive) until both rising Voltage thresholds reset limits are exceeded for 100 ms. When powering down, the drive resets immediately when either supply voltage drops below the falling voltage threshold.

Table 4-4 Power Reset Limits¹

DC VOLTAGE	THRESHOLD
+5 V	4.31/4.43v
+12 V	9.59/9.77v

4.5.3 Drive Power Dissipation

Table 4-5 lists the drive power dissipation and the corresponding currents for the various modes of operation of the Maxtor Atlas 10K V hard disk drive.

Table 4-5 Power Dissipation in Various Modes (Low-Profile Drives)

MODE OF OPERATION	TYPICAL AVERAGE CURRENT ^{1, 4} (AMP RMS UNLESS OTHERWISE NOTED)						TYPICAL AVG. POWER ⁵ (WATTS)		
	+12V			+5V					
	73.5 GB	147.1 GB	300 GB	73.5 GB	147.1 GB	300 GB	73.5 GB	147.1 GB	300 GB
Startup peak1	1.34	1.33	1.57	0.84	0.68	0.69	19.7	19.4	22.41
Idle2	0.33	0.38	0.56	0.81	0.77	0.80	8.01	8.52	10.77
Max Work	0.75	0.78	0.96	0.84	0.80	0.83	13.16	13.32	15.70

Note:

1. Current is RMS except for Startup. Startup current is the peak (> 10 ms) current required during spindle startup. Current measurements do not include power required for SCSI termination.
2. Idle mode is in effect when the drive is not reading, writing, seeking, or executing any commands. A portion of the R/W circuitry is powered down, the motor is up to speed and the Drive Ready condition exists.
3. Max Workload: Maximum workload is defined as a random read/write profile with a 50% read and 50% write distribution. The transfer length is 8 blocks (4KB), and a queue depth of 4 is used.
4. Power requirements reflect nominal values for +12V and +5V power supplies.

4.6 ACOUSTICS

Table 4-6 specifies the acoustical characteristics of the Maxtor Atlas 10K V hard disk drive. The acoustics is measured in an anechoic chamber with background noise <25 dBA.

Table 4-6 Acoustical Characteristics—Sound Power per ISO 7779

OPERATING MODE	SOUND POWER (MEAN)	SOUND POWER (MEAN + 3 SIGMA)
	MAXTOR ATLAS 10K IV 73.5/147.1/300 GB	
Idle On Track		
73.5 GB	3.16	3.38
147.1 GB	3.52	3.68
300 GB	3.67	3.91
Seeking Random		
73.5 GB	3.60	3.72
147.1 GB	3.87	4.04
300 GB	3.96	4.19

1. The specifications for idle and operating acoustic noise as per ISO Standard 7779, “Engineering Methods for Free Field Conditions Over a Reflecting PLane.” All specifications are sound power level maximum limits, A-weighted, referred to 1 picowatt as indicated in the standard. The drive will be mounted for the test in the manufacturer’s defined nominal position. The unit under test should be supported so that its bottom surface in the test orientation is less than or equal to 6 centimeters from the chamber floor, but not in contact with it. The unit should be supported with small, compliant, well-damped blocks that rest on the floor.
2. The relationship between bels and dBA for sound power is 1 bel = 10dBA.

4.6.1 Acoustic Toned Quality

A typical drive will have no discrete tones that exceeds a tone ratio of 8dB as measured per ECMA 74, 8th Edition, December 2003.

4.7 MECHANICAL

The Maxtor Atlas 10K V hard disk drives have the following mechanical characteristics:

Height¹: 1.028 in. (26.1 mm)
Width: 4.0 in. (101.6 mm)
Depth: 5.787 in. (147.0mm)
Weight (4-disk): 1.8 Ib (820 grams)

All dimensions are exclusive of any optional faceplate.

Drive mode small form factor specifications.

1. Except when shock feet are uncompressed (see of Chapter 3).

4.8 ENVIRONMENTAL CONDITIONS

Table 4-7 summarizes the environmental specifications of the Maxtor Atlas 10K V hard disk drive.

Table 4-7 Environmental Specifications

PARAMETER	OPERATING	NON-OPERATING
Temperature (Non-condensing)	5° to 55°C (41° to 131°F)	-40° to 70°C (-40° to 158°F)
Maximum HDA Temperature	60°C	N/A
Temperature Gradient (Non-condensing)	20°C/hr maximum	30°C/hr maximum
Humidity ¹ Maximum Wet Bulb Temperature	5% to 95% RH 32°C (89.6°F)	5% to 95% RH 46°C (114°F)
Humidity Gradient	30% per hour	30%per hour
Altitude ²	-300 m to 3,048 m (-1,000 to 10,000 ft.)	-300 m to 12,000 m (-1,000 to 40,000 ft.)

1. No condensation.
2. Altitude is relative to sea level. 1,000 feet = 305 meters.

1.

4.9 ELECTROMAGNETIC CONDITIONS

4.9.1 EMI/RFI Susceptibility

3 Volts/meter 80% modulated from 80 to 1000 MHz.

4.9.2 ESD

Drive must function with no data loss or component damage with air discharges of 1 to 15 KV, and contact discharges of 2 to 8 KV in both positive and negative voltages.

4.9.3 Sensitivity to Magnetic Fields

The Maxtor Atlas family of drives meet 3A @ 230V / 50 Hz specification.

4.10 SHOCK AND VIBRATION

The Maxtor Atlas 10K V hard disk drive can withstand levels of shock and vibration applied to any of its three mutually perpendicular axes, or principal base axis, as specified in Table 4-8. A functioning drive can be subjected to specified operating levels of shock and vibration. When a drive has been subjected to specified non-operating levels of shock and vibration, with power to the drive off, there will be no change in performance at power on.

When packed in its 1-pack shipping container, the Maxtor Atlas 10K V drive can withstand a drop from 42 inches onto a concrete surface on any of its surfaces, six edges, or three corners. The 20-pack shipping container can withstand a drop from 36 inches onto a concrete surface on any of its surfaces, six edges, or three corners.

Table 4-8 Shock and Vibration Specifications

PARAMETER	OPERATING	NON-OPERATING
Shock linear Half Sine 2 msec ¹ Read/Write	63G/30G ¹	250 G ¹
Shock linear Half Sine 11 msec ¹	15 G ¹	100 G ¹
Shock rotational (rad/s ²) ¹ , 2 msec	7000	25,000
Vibration Swept ¹	1.5G 0-p 5-500 Hz	2.0G 0-p 5-500 Hz
Vibration Random (G ² /Hz)	0.008 10—300 Hz 0.0012 300—500 Hz	0.014 10—300 Hz 0.01 300—500 Hz

Note: 1. At ambient temperature and no unrecoverable errors.

4.11 RELIABILITY

Component Life: 5 years

Preventive Maintenance (PM): Not required

Contact Start/Stop: 50,000 cycles at ambient,
35,000 at Environments

Annualized Failure Rate (AFR) No greater than 0.58 percent

Drive reliability is closely related to the temperatures the drive is exposed to. The AFR is based on an operational design temperature ambient of 86° F (30° C)

4.12 DISK ERRORS

Table 4-9 provides the error rates for the Maxtor Atlas 10K V hard disk drive.

Table 4-9 Error Rates

ERROR TYPE	MAXIMUM NUMBER OF ERRORS
Recovered read errors ¹	<10 events per 10 ¹² bits read
Uncorrectable read errors ²	1 event per 10 ¹⁵ bits read

1. Recovered read errors are errors which require retries for data correction. Errors corrected by ECC on-the-fly are not considered recovered read errors. Read on arrival is disabled to meet this specification.
2. Uncorrectable read errors are those that are not correctable using an Error Correcting Code (ECC), or retries within the retry limits specified in the mode pages. The drive terminates the command either when a repeating error pattern occurs, or after both the specified number of retries and application of triple burst error correction fail to correct the error.
3. A seek error occurs when the actuator fails to reach or remain on the requested cylinder, and/or the drive requires the execution of the full calibration routine to locate the requested cylinder.

Chapter 5

SCSI Description

This chapter contains an overview of SCSI command and status processing and a detailed description of the commands supported by the disk drives. The SCSI command system enables the initiator to instruct the drive to perform specific functions.

In this manual, unless otherwise stated, numerical values are given in decimal. Hexadecimal numbers, such as opcodes, are always given with an “h” following, as in 5Ah except when entire data tables are in hexadecimal.

5.1 Overview of the SCSI Command Descriptions

The disk drives support the SCSI-3 commands listed in Table 5-1. The command categories are *sequential*, *normal*, and *immediate*.

Immediate commands are processed when received by the drive. In most cases, these commands do not require drive resources, do not change the state of the drive, and bypass the command queue (unless the command is tagged).

Sequential commands execute in the order that they arrive (unless accompanied by a HEAD OF QUEUE tag) and execute to completion before a subsequent command is activated.

Normal commands are allowed to execute concurrently (with the restriction that the drive executes overlapping writes in the order received). Normal commands are usually I/O commands

Table 5-1 Supported SCSI Commands

Command	Operation Code	Category
CHANGE DEFINITION	40h	Sequential
FORMAT UNIT	04h	Sequential
INQUIRY	12h	Immediate
LOG SELECT	4Ch	Sequential
LOG SENSE	4Dh	Sequential
MODE SELECT (6)	15h	Sequential
MODE SELECT (10)	55h	Sequential
MODE SENSE (6)	1Ah	Sequential
MODE SENSE (10)	5Ah	Sequential
PERSISTENT RESERVATION IN	5Eh	Sequential

Command	Operation Code	Category
PERSISTENT RESERVATION OUT	5Fh	Sequential
READ (6)	08h	Normal
READ (10)	28h	Normal
READ BUFFER	3Ch	Sequential
READ CAPACITY	25h	Immediate
READ DEFECT DATA (10)	37h	Sequential
READ DEFECT DATA (12)	B7h	Sequential
READ LONG	3Eh	Sequential
REASSIGN BLOCKS	07h	Sequential
RECEIVE DIAGNOSTIC RESULTS	1Ch	Sequential
RELEASE	17h	Sequential
RELEASE (10)	57h	Sequential
REPORT LUNS	A0h	Sequential
REPORT DEVICE IDENTIFIER	A3h	Sequential
REQUEST SENSE	03h	Immediate
RESERVE	16h	Sequential
RESERVE (10)	56h	Sequential
REZERO UNIT	01h	Sequential
SEEK (6)	0Bh	Sequential
SEEK (10)	2Bh	Sequential
SEND DIAGNOSTIC	1Dh	Sequential
SET DEVICE IDENTIFIER	A4h	Sequential
START STOP UNIT	1Bh	Sequential
SYNCHRONIZE CACHE	35h	Sequential
TEST UNIT READY	00h	Immediate
VERIFY	2Fh	Normal
WRITE (6)	0Ah	Normal
WRITE (10)	2Ah	Normal
WRITE AND VERIFY	2Eh	Normal
WRITE BUFFER	3Bh	Sequential
WRITE LONG	3Fh	Sequential
WRITE SAME	41h	Sequential

-
1. Relative Addressing is not supported by the drive. Therefore, in all I/O commands, the RelAdr bit must be 0.
 2. RESERVE and RELEASE are supported, as are third-party reservations. Extent reservations are not supported.
 3. The RECEIVE DIAGNOSTIC RESULTS and SEND DIAGNOSTIC DATA commands implement vendor-unique pages to test the drive during the manufacturing process. It is recommended that initiators specify only the non-page format variants of these commands (PF=0), except for page 0x40.

5.2 Command Descriptor Block

An initiator communicates with the drive by sending a 6-, 10-, or 12-byte Command Descriptor Block (CDB) that contains the parameters for the specific command. The SCSI command's operation code is always the first byte in the CDB and a control field is the last byte. For some commands, the CDB is accompanied by a list of parameters sent during the data-out buffer transfer. Table 5-2 shows the format of a typical six-byte CDB.

Table 5-2 Typical (6-Byte) Command Descriptor Block-Data Format

Bit Byte	7	6	5	4	3	2	1	0	
0	Operation Code								
1	Reserved			(MSB)	Logical Block Address				(LBA)
2 - 3	Logical Block Address (LBA)								
4	Transfer Length, Parameter List Length, or Allocation Length								
5	Control								

Note: Reserved fields in CDBs and Parameters have a value of 0; Reserved fields in states and other parameters sent to an initiator are set to 0.

Table 5-3 contains a description of the CDB fields.

Table 5-4 describes the data format for a typical CDB.

Table 5-3 Command Descriptor Block-Field Descriptions

Field	Description																		
Operation Code	The first byte of a SCSI CDB contains an operation code. The operation code of the CDB has a <i>Group Code</i> field (bits 7-5) and a <i>Command Code</i> field (bits 4-0). The 3-bit Group Code field provides for eight groups of command codes. The 5-bit Command Code field provides for 32 command codes in each group. A total, therefore, of 256 possible operation codes exist. Operation codes are defined in the SCSI command standards. The group code for CDBs specified therein correspond to the length of the command descriptor as below:																		
	<table border="1"> <thead> <tr> <th>Group Code</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>6-byte commands</td> </tr> <tr> <td>1</td> <td>10-byte commands</td> </tr> <tr> <td>2</td> <td>10-byte commands</td> </tr> <tr> <td>3</td> <td>Reserved</td> </tr> <tr> <td>4</td> <td>16-byte commands</td> </tr> <tr> <td>5</td> <td>12-byte commands</td> </tr> <tr> <td>6</td> <td>Vendor specific</td> </tr> <tr> <td>7</td> <td>Vendor specific</td> </tr> </tbody> </table>	Group Code	Meaning	0	6-byte commands	1	10-byte commands	2	10-byte commands	3	Reserved	4	16-byte commands	5	12-byte commands	6	Vendor specific	7	Vendor specific
Group Code	Meaning																		
0	6-byte commands																		
1	10-byte commands																		
2	10-byte commands																		
3	Reserved																		
4	16-byte commands																		
5	12-byte commands																		
6	Vendor specific																		
7	Vendor specific																		
	The operation code specifies the command being requested. The list of supported SCSI commands and their operation codes are contained in Table 5-1.																		

Field	Description
Logical Block Address	<p>The 6-byte READ, SEEK, and WRITE Command Descriptor Blocks contain a 21-bit Logical Block Address. The 10-, 12-, and 16- Command Descriptor Blocks contain a 32-bit Logical Block Address.</p> <p>Commands that require additional parameter data specify the length of the Logical Block Address that is needed. See the specific command descriptions for more detailed information.</p> <p>Relative Addressing indicates a technique used to determine the next Logical Block Address to be operated on. The drive does not support Relative Addressing, it defaults to a value of 0, which specifies that the Logical Block Address specifies the first logical block of a range of logical blocks to be operated on by the command.</p>
Transfer Length	<p>The transfer length field normally specifies the number of sectors to be transferred between the initiator and the drive. For several commands, the transfer length indicates the number of bytes (not sectors) to be sent. For these commands, this field may be identified by a different name</p> <p>Commands that use <i>one byte</i> for the transfer length value allow up to 256 sectors of data to be transferred by one command. A transfer length value of 0 indicates that 256 sectors are to be sent. Transfer length values of 1 through 255 indicate the number of sectors to be transferred.</p> <p>Commands that use <i>multiple bytes</i> for the transfer length value function differently. A transfer length value of 0 indicates that no data transfer is to occur. Transfer length values of 1 or greater indicate the number of sectors to be transferred.</p>
Parameter List Length	<p>The Parameter List Length is used to specify the number of bytes sent during the data-out buffer transfer. This field is typically used for parameters that are sent to a drive (for example, mode, diagnostic, and log parameters). A parameter list length of 0 indicates that no data is to be transferred.</p>
Allocation Length	<p>The Allocation Length field specifies the maximum number of bytes that the initiator has allocated for returned data. The Allocation Length is used to limit the amount of data returned to the initiator.</p> <p>An Allocation Length of 0 indicates that no data is to be transferred from the drive to the initiator. The drive terminates the data-in buffer transfer when the specified number of bytes have been transferred to the initiator or when all available data has been transferred, whichever is less.</p>
Control Field	<p>The Control Field is the last byte of every Command Descriptor Block; its format is shown in Figure 5-2 and described in Table 5-3.</p>

Table 5-4 Typical (6-Byte) Command Descriptor Block - Data Format

Bit	7	6	5	4	3	2	1	0
Byte								
5	Vendor Specific		Reserved			NACA	Flag	Link

Table 5-5 contains a description of the CDB control field.

Table 5-5 Command Descriptor Block Control Field-Field Descriptions

Field	Description
Vendor Specific Bits	These bits must be 0.
NACA	Normal Auto-Contingent Allegiance - This bit must be zero to indicate that SCSI-2 Contingent Allegiance rules apply.
Link Bit	<p>A Link bit set to one signals that the initiator requests continuation of a task (I/O Process) across two or more SCSI commands. If the Link bit is one and the flag bit is zero, and the command completes successfully, the drive will continue the task and return a status of INTERMEDIATE and a service response of Linked Command Complete.</p> <p>If the Link bit and the Flag bit of the Control word are both set to one, and the drive completes a command with a status of INTERMEDIATE, the drive will return a service response of Linked Command Complete (with Flag). Refer to Section 5.4 for Linked Commands description.</p>
Flag Bit	The Flag bit is used in conjunction with the Link Bit to notify the initiator in an expedient manner that a command has been completed. A flag bit set to 1 is valid only when the Link Bit is set to 1.

5.3 Status/Error Reporting

SCSI message-level errors are communicated by messages that are defined specifically for that purpose. SCSI command-level errors are communicated by a status that is returned by the drive during the STATUS phase. This phase occurs at the end of each command, unless the command is terminated by one of the following events:

- ABORT TASK SET message
- ABORT TASK message
- TARGET RESET message
- CLEAR QUEUE message
- Unexpected disconnect

The status code is contained in bits 1 through 5 of the status byte. Bits 0, 6, and 7 are reserved. Table 5-6 describes the status codes returned by the drive.

Table 5-6 Status Codes

Status	Definition	Meaning
00h	GOOD	The drive successfully completed the command.
02h	CHECK CONDITION	An Auto Contingent Allegiance (ACA) condition occurred.
08h	BUSY	The drive cannot service the command at the moment, and its Command Descriptor Block has been discarded. The initiator can retry the command at a later time. This status is returned when: <ul style="list-style-type: none">• A non-tagged command is received and the logical unit's command queue is full (all internal command buffers are in use).• A disconnect privilege was not granted in the IDENTIFY message of a queue-tagged I/O process (Parallel SCSI).• A disconnect privilege was not granted in the IDENTIFY message of a non-tagged I/O process and a command from another initiator is currently active (Parallel SCSI)• A command is received while an auto-contingent allegiance condition exists for another initiator.
10h	INTERMEDIATE	This status is returned for every command (except the last) in a series of linked commands that was successfully completed. However, if the command is terminated with other than GOOD status (such as CHECK CONDITION, RESERVATION CONFLICT, OR BUSY), the INTERMEDIATE status is not returned and the series of linked commands and the task is ended.
18h	RESERVATION CONFLICT	Another initiator has reserved the drive. (This status is never returned for INQUIRY or REQUEST SENSE commands.)
28h	TASK SET FULL	The drive cannot service the command at the moment, and its Command Descriptor Block has been discarded. (Returned for a tagged command when all of the drive's internal command buffers are in use, or when a host sends a tagged command while an Auto Contingent Allegiance condition is pending for that initiator)

5.4 Linked Commands

An I/O Process (task) may contain multiple commands that are linked together. The initiator communicates this condition of linked (or unlinked) commands by setting (or clearing) the Link bit of the Command Descriptor Block's control word. A linked

command is one in which the Link bit in the Command Descriptor Block is set. After successful completion of a linked command the drive sends an INTERMEDIATE status, followed by a LINKED COMMAND COMPLETE message. If the Flag bit was set in the Command Descriptor Block, the drive sends an INTERMEDIATE status, followed by a LINKED COMMAND COMPLETE (WITH FLAG) message. The drive then switches the bus to the command phase in order to receive the next command in the linked chain.

All commands in a linked chain are addressed to the same nexus and are part of a single task (I/O process). The drive defers any commands that are not part of the linked chain until the chain is complete. The last command in the chain has the Link bit cleared.

Note: Relative addressing is not supported by the Atlas 10K disk drive.

5.5 DATA Transfer Command Components

Many of the SCSI commands cause data to be transferred between the initiator and the drive. The content and characteristics of this data are command-dependent. Table 5-8 lists the information transmitted for all of the commands.

The “Length in CDB” column of Table 5-8 identifies the Command Descriptor Block field used by the drive to determine how much command-related data are to be transferred. The units (bytes or logical blocks) for the different Length fields are implied by the Length Field Name as shown in Table 5-7:

Table 5-7 Length Fields

Field Name	Units Implied
Allocation Length	Bytes of data the drive is allowed to send to the initiator
Parameter List Length	Bytes of data the initiator has available for the drive
Transfer Length	Logical data sectors the initiator wants transferred or verified
Byte Transfer Length	Bytes of data the initiator wants transferred

The DATA OUT column in Table 5-8 lists the information passed to the drive by the initiator as part of the command. The DATA IN column lists the information sent to the initiator by the drive.

Numbers in parentheses after an item indicate the item’s length in bytes. In some cases, additional length information is communicated during the DATA phase. For example, a FORMAT UNIT Defect List Header contains a Defect List Length field that contains the total length of the Defect Descriptors that follow the Defect List Header. Table 5-8 does not include these cases.

Table 5-8 DATA-Phase Command Contents

Command	Length in CDB	Data Out (To Drive)	Data In (To Initiator)
CHANGE DEFINITION	0	--	--
FORMAT UNIT	0	Defect List Header Initialization Pattern (6-8) Defect Descriptors	
INQUIRY	Allocation	---	Standard Inquiry or a Vital Product Data page
LOG SELECT	Parameter List (must be 0)	--	--
LOG SENSE	Allocation	---	Log Page
MODE SELECT	Parameter List	Mode Parameter Header (4 Block Descriptor (8 Page(s))	--
MODE SENSE	Allocation	---	Block Descriptor (8)
PERSIST. RES. IN			
PERSIST. RES. OUT			
READ (6) (10)	Transfer	---	Data

Command	Length in CDB	Data Out (To Drive)	Data In (To Initiator)
READ BUFFER	Allocation	---	READ BUFFER Header (4) Mode-zero Buffer (512) or Section of Drive's DRAM or READ BUFFER Desc. (4)
READ CAPACITY	Allocation	---	READ CAPACITY data (8)
READ DEFECT DATA	Allocation	---	<ul style="list-style-type: none"> • Defect List (Hdr) (4) • Defect Descriptors
READ LONG	Byte Transfer (Must be 578)	---	<ul style="list-style-type: none"> • Data (512) • LBA Tag (2) • EDC (2) • ECC (60) • Fill (2)
REASSIGN BLOCKS	0	Defect List Header(4)	Defect Descriptors
RECEIVE DIAGNOSTIC RESULTS	Allocation	---	Diagnostic Page
RELEASE	0	---	---
REPORT DEVICE IDENTIFIER			
REQUEST SENSE	Allocation	---	Sense Data (18)
RESERVE	0 (Extent List Option not supported)	---	---
REZERO UNIT	0	---	---
Command	Length in CDB	Data Out (To Drive)	Data In (To Initiator)
SEEK (6)(10)	0	---	---
SEND DIAGNOS- TIC	Parameter List	Diagnostic Page	---
SET DEVICE IDENTIFIER			
START STOP UNIT	0	---	---
SYNCHRONIZE CACHE	0	---	---
TEST UNIT READY	0	---	---
VERIFY	Transfer	Data	---
WRITE (6)(10)	Transfer	Data	---
WRITE AND VERI- FY	Transfer	Data	---
WRITE BUFFER	Parameter List	Zeros (4) Mode-zero buffer (512) or Data to put into DRAM or Microcode image or Microcode image (successive 16 KB pieces)	---

Command	Length in CDB	Data Out (To Drive)	Data In (To Initiator)
WRITE LONG	Byte Transfer (must be 578)	<ul style="list-style-type: none"> • Data (512) • LBA Tag (2) • EDC (2) • ECC (60) • Fill (2) 	--
WRITE SAME	0	Data (1 logical sector)	

5.6 SCSI COMMAND DESCRIPTIONS

The SCSI command descriptions that follow this page contain detailed information about the SCSI commands that are supported by the drive. Each description provides a Data Format and Field Descriptions for the Command Descriptor Block for the described command.

The commands are presented in alphabetic order, and each command starts on a new, odd-numbered page.

Common Fields

Several fields that are common to many commands are described here, rather than being repeated throughout the descriptions. These fields include:

Reserved – Reserved bits, fields, bytes, and code values are set aside for future standardization and must be set to 0. If the drive receives a command that contains non-0 bits in a reserved field or a reserved code value, the command is terminated with a CHECK CONDITION status and the sense key set to ILLEGAL REQUEST. However, there are some fields that are not checked for compatibility with older SCSI initiators.

Control – The Link Bit and Flag Bit are supported.

RelAdr – Not supported; **must be 0**.

5.7 CHANGE OPERATING DEFINITION (40h)

The CHANGE DEFINITION command shown in Table 5-9 and Table 5-10 set the SCSI compliance for disk drives with parallel SCSI interfaces to one of four different levels: SCSI-1, SCSI-1/CCS, SCSI-2 or SCSI-3.

Table 5-9 *CHANGE DEFINITION Command Descriptor Block-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (40h)							
1	Reserved							
2	Reserved							Save
3	Rsv'd	New Operating Definition (0, 1, 2, 3, 4, 5 and 6 are legal values for Parallel SCSI Disk Drives)						
4 - 7	Reserved							
8	Parameter Data Length							
9	Control							

Table 5-10 *CHANGE DEFINITION-Field Descriptions*

Field	Description
Save	Save Parameters. When this field is set to 1, the new operating definition is saved in the drive's non-volatile memory.
New Operating Definition	The values that can be supplied in the field are listed below: 0 Current Definition 1 SCSI-1 2 SCSI-1/CCS (Common Command Set) 3 SCSI-2 4 SCSI-3 5 SCSI-3 U160 6 SCSI-3 U320
Parameter Data Length	Length, in bytes, of the parameter data. Must be 0.

5.8 FORMAT UNIT Command (04h)

The FORMAT UNIT command formats the disk's storage media into initiator-addressable logical blocks according to initiator-defined options. This command ensures that the disk storage media is formatted so that all data sectors are accessible. Any data residing on the disk before this command is invoked is lost. This command repairs damage left by previous WRITE LONG commands. The WRITE LONG command allows the initiator to deliberately corrupt a sector's ECC or EDC.

The FORMAT UNIT command updates the defective sector list, referred to as the Grown Defect List, that is maintained by the drive. As described below, the drive gives the initiator considerable control over this list. If certification is enabled, all initiator-addressable logical blocks are verified, and bad sectors are reassigned and added to the Grown Defect List.

Note: During FORMAT UNIT command processing, the drive ignores the Read/Write AWRE and ARRE bits (from Error Recovery Mode page) and does auto-revectoring as needed.

The drive is physically formatted when it leaves the manufacturing facility. Therefore, it is not necessary to reformat the drive before using it. If the logical sector size of the drive is changed using the Block Descriptor of the MODE SELECT Command, it is recommended (but not necessary) to reformat the drive.

The Immed bit that can be included in the FORMAT UNIT command parameter list allows the initiator to control whether the drive returns completion status either after fetching and validating the Command Descriptor Block and parameter list or after the FORMAT UNIT command completes. The FORMAT UNIT command parameters are described in Table 5-11 and Table 5-12

The time required for the FORMAT UNIT command to format the disk primarily depends on the capacity of the drive. The number of defects detected and the number of defects already in the Primary and Grown Defect Lists also influence the time required to format the drive. If the Disable Certification bit is set to 1, formatting time is reduced.

Defective sectors on the drive are managed through two lists: the Primary Defect List and the Grown Defect List. The Primary Defect List is created when the drive is manufactured and is the drive's initial defect list. The Primary Defect List is not affected by the FORMAT UNIT command. Sectors listed in it are revectoring by the drive. The Grown Defect List contains a list of the sectors that have gone bad since the drive's primary list was generated.

Defects are communicated in a data structure referred to as a Defect Descriptor. The SCSI specification defines several Defect Descriptor formats. The FORMAT UNIT command recognizes defect descriptors in Block format, Bytes From Index format, and Physical Sector format. (There is one exception: a value of FFFFFFFFh in either the defect Bytes From Index field or the Defective Sector Number fields is ignored).

Table 5-11 *FORMAT UNIT Command Descriptor Block-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (04h)							
1	Reserved			Fmt Data	Cmp List	Defect List Format		
2	Vendor-Specific							
3 - 4	Interleave							
5	Control							

Table 5-12 *FORMAT UNIT Command-Field Descriptions*

Field	Description
FmtData	<p>Format Data. When the Format Data value is 1, it indicates that a data-out buffer transfer occurs as part of the command execution. The FORMAT UNIT Parameter list (consisting of a Defect List Header and, optionally, a number of Block Format Defect Descriptors) is passed to the device during this phase.</p> <p>When the Format Data value is 0, the data-out buffer transfer does not occur.</p>
CmpLst	<p>Complete List. When the Complete List value is 1, the drive deletes its current Grown Defect List and starts a new one, containing the Logical Block Numbers listed in the defect list supplied by the initiator during this format operation are added to the list, creating a new Grown Defect List.</p> <p>When the Complete List value is 0, the drive adds initiator-supplied and newly found defective Logical Block Numbers to the existing Grown Defective List.</p>
Defect List Format	<p>The Defect List Format value specifies the defect descriptor passed by the initiator to the drive when the Format Data value is 1. Acceptable values are:</p> <ul style="list-style-type: none"> 000b (Block [<i>or Sector</i>] Format) 100b (Index Format) 101b (Physical Sector Format)
Vendor-Specific	Not supported. Must be 0.
Interleave	Not supported. Ignored by the drive.

5.8.1 Five Forms of FORMAT UNIT Commands

Five different forms of the FORMAT UNIT command are supported through different combinations of the Format Data bit, Complete List bit, and the information in the Defect List Length field. (Refer to *FORMAT UNIT Defect Header List* for a description of the Defect List Length field.) The different command forms give the initiator control over the contents of the Grown Defect List. Table 5-13 describes the contents of the Grown Defect List after FORMAT UNIT command execution.

Table 5-13 *FORMAT UNIT Command Supported Options*

FmtDat	CmpLst	Defect List Length	Contents of Grown Defect List after FORMAT UNIT Command Execution
0	0	N/A	All sectors found to be bad, but not listed in the Primary Defect List or Grown Defect List, are added to the Grown Defect List.
1	0	0	Same as above.
1	1	0	The existing Grown Defect List is discarded. A new Grown Defect List is generated, containing All sectors found to be bad, but not listed in the Primary Defect List.
1	0	>0	The Grown Defect List contains: Original Grown Defect List All sectors found to be bad, but not listed in the Primary Defect List. A list of sectors supplied by the initiator. These sectors are passed in defect descriptors in the data-out buffer transfer that occurs as part of the FORMAT UNIT command.
1	1	>0	Same as the case above, except that the current Grown Defect List is discarded before formatting begins.

5.8.2 FORMAT UNIT Parameter List

The FORMAT UNIT Parameter List (Table 5-14) consists of a Defect List Header (Table 5-15) followed by one or more Defect Descriptors. Descriptors are either four bytes or eight bytes in length.

Table 5-14 *FORMAT UNIT Parameter List-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0 - 3	Defect List Header							
	Initialization Pattern Descriptor (if any)							
	Defect Descriptors (if any)							

5.8.2.1 **FORMAT UNIT Defect List Header**

The FORMAT UNIT Defect List Header (Table 5-15) provides several optional format control bits to give the initiator more control over the defect lists. Table 5-16 provides descriptions of the data fields in the Defect List Header

Table 5-15 .FORMAT UNIT Defect List Header-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0	Reserved							
1	FOV	DPRY	DCRT	STPF	IP	DSP	Immed	VS
2-3	Defect List Length							

Table 5-16 FORMAT UNIT Defect List Header-Field Descriptions

Name	Description
FOV	The Format Options Valid bit indicates that the remaining option bits in this byte are valid. If this bit is not set, the remaining bits, except Immed, are ignored.
DPRY	The Disable Primary bit specifies whether the drive should ignore the Primary Defect List during the format to identify defective areas on the media. The drive's default is 0, indicating that the drive should replace sectors found in the Primary Defect List during the format.
DCRT	The Disable Certification bit specifies whether the drive should certify the media during the format. The drive's default is 0, indicating that the drive should certify the media.
STPF	The Stop Format bit specifies the error conditions under which the drive is to terminate the format. The state of the bit is ignored and the drive acts as if this bit is set. If either the Grown Defect List or the Primary Defect List is not found, the format operation terminates with a CHECK CONDITION status and a sense key of MEDIUM ERROR
IP	The Initialization Pattern bit signals the drive that the FORMAT UNIT Parameter List contains initialization pattern information. An Initialization Pattern bit of 1 indicates that an initialization pattern descriptor (Table 5-17) is included in the FORMAT UNIT parameter list immediately following the Defect List Header. An Initialization Pattern bit of 0 indicates that an initialization pattern descriptor (Table 5-17) is not included in the FORMAT UNIT parameter list, and the drive will use its default initialization pattern of all zeros.
DSP	The Disable Saving Parameter prohibits the drive from preserving MODE SELECT parameters received while the Format is in progress. This bit is ignored, and the drive acts as if it were set.
Immed	When the Immediate bit is 0, the drive returns a status after the format operation completes. If the disconnect privilege is granted in the accompanying IDENTIFY message, the drive disconnects from the bus after fetching and validating the Command Descriptor Block and the FORMAT UNIT Parameter List. If the disconnect privilege is not granted, the drive stays connected to the bus during the entire FORMAT UNIT command. When the Immediate bit is set to 1, the drive returns a status after the Command Descriptor Block and FORMAT UNIT Parameter List have been fetched and validated. In this case, the drive ignores the IDENTIFY message's disconnect privilege bit and remains connected to the bus while it validates the Command Descriptor Block and Parameter List. It returns the completion status before it disconnects. The drive then proceeds with the format.
VS	Vendor-specific. Must be 0.
Defect List Length	Gives the length in bytes of the Defect Descriptors that follow. A value of 0 is valid and means that no Defect Descriptors follow. A CHECK CONDITION status is returned if this value is not a multiple of the defect descriptor size. The Defect List Length is equal to four times the number of defect descriptors if Block format is used, or eight times if Bytes From Index format or Physical Sector format is used.

5.8.2.2 FORMAT UNIT Defect Descriptor-Block Format

Each descriptor specifies a 4-byte defective sector address of the sector that contains the defect as shown in Table 5-17.

Table 5-17 *FORMAT UNIT Defect Descriptor-Block Format*

Bit Byte	7	6	5	4	3	2	1	0
0 - 3	Defective Block Address							

5.8.2.3 FORMAT UNIT Defect Descriptor — Physical Sector and Bytes From Index Format

The Physical Sector defect descriptor specifies the location of a defect that is the length of a sector. The Bytes From Index defect descriptor specifies the location of a defect that is no more than eight bytes long.

Each descriptor is comprised of the bytes shown in Table 5-18. The Cylinder Number of Defect is the most significant part of the address, and the Defective Sector Number or Defect Bytes From Index is the least significant part of the address. A value of FFFFFFFFh in bytes 4 to 7 is ignored

Table 5-18 *FORMAT UNIT Defect Descriptor- Physical Sector and Bytes From Index Format*

Bit Byte	7	6	5	4	3	2	1	0
0 - 2	Cylinder Number of Defect							
3	Head Number of Defect							
4 - 7	Defective Sector Number or Defect Bytes from Index							

5.8.2.4 FORMAT UNIT Initialization Pattern Descriptor

The Initialization Pattern option specifies that the logical blocks on the drive medium will contain a specified initialization pattern. The Initialization Pattern descriptor is sent to the drive as part of the FORMAT UNIT parameter list. Table 5-19 describes the Initialization Pattern descriptors. The contents of the descriptor fields are described in Table 5-20 and the types in Table 5-21.

Table 5-19 *FORMAT UNIT Initialization Pattern Descriptor-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0	IP Modifier = 0		Reserved					
1	IP Type							
2 - 3	IP Length							
4 - n	IP							

Table 5-20 *FORMAT UNIT Initialization Pattern Descriptor-Field Descriptions*

Name	Description
IP Modifier	The Initialization Pattern Modifier must be 0.
IP Type	The Initialization Pattern Type field (Table 5-21) indicates the type of pattern the drive uses to initialize each logical sector within the initiator-accessible portion of the medium. All bytes within a logical sector are written with the initialization pattern.
IP Length	The Initialization Pattern Length field indicates the number of bytes contained in the Initialization Pattern. The valid lengths (when IP Type = 1) are two or four bytes.
IP	Initialization Pattern.

Table 5-21 *FORMAT UNIT Initialization Pattern Type*

Initialization Pattern Type	Description
00h	Use default pattern. ¹
01h	Repeat the initialization pattern as required to fill the logical sector. ²
02h – FFh ³	Reserved.

NOTES:

1. If the Initialization Pattern Length is not 0, the drive terminates the command with CHECK CONDITION status. The sense key is set to ILLEGAL REQUEST, and the additional sense code is set to INVALID FIELD IN PARAMETER LIST.
2. If the Initialization Pattern Length is 0, the drive terminates the command with CHECK CONDITION status. The sense key is set to ILLEGAL REQUEST, and the additional sense code is set to INVALID FIELD IN PARAMETER LIST.
3. If the Initialization Pattern Type is not 0 or 1, the drive terminates the command with a CHECK CONDITION status. The Sense Key is set to ILLEGAL REQUEST, and the Additional Sense Code is set to INVALID FIELD IN PARAMETER LIST.

5.9 INQUIRY Command (12h)

The INQUIRY command allows the initiator to determine the kind of SCSI devices attached to its SCSI bus. It causes a device that is attached to a SCSI bus to return information about itself. The drive identifies itself as a Direct Access Storage Device that implements the applicable interfacing protocol. The drive does not need to access its storage medium to respond to the inquiry. The INQUIRY commands are described in Table 5-22 and Table 5-23.

The drive can provide two categories of data in response to an INQUIRY command: Standard Inquiry Data and Vital Product Data. Standard Inquiry Data contains basic data about the drive, and Vital Product Data comprises several pages of additional data. Each Vital Product Data page requires a separate INQUIRY command from the initiator.

An INQUIRY command is not affected by, nor does it clear, a Unit Attention condition.

Table 5-22 *INQUIRY Command Descriptor Block-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (12h)							
1	Reserved						CmdDt	EVPD
2	Page Code or Operation Code							
3	Reserved							
4	Allocation Length							
5	Control							

Table 5-23 *INQUIRY Command Descriptor Block-Field Descriptions*

Field	Description
CmdDt	Command Support Data. If CmdDt = 0 and EVPD (see below) = 0, the drive returns the Standard Inquiry Data. If CmdDt = 1 with EVPD = 0, the drive returns the Command Data specified by Page Code/Operation.
EVPD	Enable Vital Product Data. If EVPD = 0 and CmdDt (see above) = 0, the drive returns the Standard Inquiry Data. If EVPD = 1 and CmdDt = 0, the drive returns the Vital Product Data Page specified by Page Code/Operation Code.
Page Code or Operation Code	Specifies the Vital Product Data Page or Command Support Data which is to be returned by the drive when EVPD is set. Specifies the SCSI Operation Code for command support data to be returned by the drive when CmdDt is set. A CHECK CONDITION status is returned if this field specifies an unsupported Page or Operation Code or if both EVPD and CmdDt are set. Table 5-27 lists the Page Codes for the Vital Product Pages supported by the drive. Table 5-37 lists the operation codes for the CmdDt information supported by the drive.
Allocation Length	Specifies the number of bytes of inquiry information the drive is allowed to return to the initiator during the command's data-in buffer transfer. Error status is not returned if the value in this field truncates the requested information.

5.9.1 Standard Inquiry Data Page

The Standard Inquiry Data Page is returned by the drive in response to the INQUIRY command if EVPD = 0. The data format is described in Table 5-24 and the fields are described in Table 5-25.

Table 5-24 *Standard Inquiry Data Page-Data Format*

Byte	Bit	7	6	5	4	3	2	1	0
0		Peripheral Qualifier = 0			Peripheral Device Type = 0 (Direct Access Device)				
1		RMB = 0	Device Type Modifier = 0						
2		ISO Version = 0		ECMA Version = 0			ANSI Version = 3 (SCSI-3)		
3		AERC	TmTsk	Norm ACA	Rsv'd	Response Data Format = 2			
4		Additional Length = 5Bh							
5		Reserved							
6		Rsv'd	Rsv'd	Port	MultiP	Mchngr	Obsolete	Obsolete	Addr16
7		RelAdr	Obsolete	Wbus16	Sync	Linked	Obsolete	Cmd Que	SftRe
8 - 15		Vendor Identification "MAXTOR"							
16 - 31		Product Identification							
32 - 35		Product Revision Level							
36 - 47		Drive Serial Number							
48 - 51		Reserved							
52 - 53		SCSI Hardware Revision Number							
54		Disk Controller Hardware Revision Number							
55		Electronics Pass Number							
56		Reserved				Clocking		QAS	IUS
57 - 95		Reserved							

Table 5-25 Standard Inquiry Data Page-Field Descriptions

Field Name	Value	Description										
Peripheral Qualifier	0	Non-zero if initiator selects an invalid logical unit.										
Peripheral Device	0	0 indicates that this is a Direct Access Device.										
ANSI Version	3	ANSI SCSI Level 3 (SCSI-3) is supported.										
AERC	0	Asynchronous Event Reporting is not supported.										
NormACA	0	Does not support setting NACA in CDB Control word.										
Port	0	Only used when MultiP = 1.										
MultiP	0	This field set to 1 if it is a multiport device.										
Mchngr	0	Not embedded in or attached to a medium changer										
Addr16	0/1	16-bit wide SCSI address										
TrmTSK	0	TERMINATE TASK Function is not supported										
Response Data Format	2	This Standard Inquiry Data is in the format specified in the International Standard.										
RelAdr	0	Relative Addressing is not supported.										
WBus 16	0 or 1	The WBus bit is 1 if the drive supports 16-bit data transfer. The bit is 0 if the drive only supports 8-bit transfer.										
Sync	1	The drive supports Synchronous Data Transfers.										
Linked	1	Linked Commands are supported.										
CmdQue	1	The drive supports Tagged Command Queuing										
SftRe	0	The drive implements the hard reset option in response to assertion of the SCSI Bus reset line.										
VS	0	Vendor Specific										
Product Identification	The value in this field is:	ATLAS10K5_73WLS, ATLAS10K5_147WLS, ATLAS10K5_300WLS ATLAS10K5_73SCA, ATLAS10K5_147SCA, ATLAS10K5_300SCA										
Clocking	11	The contents of this field define the setting of the double-edge clocking option (See Chapter 6). Note that the clocking field does not apply to asynchronous transfers of data. Options available are:										
		<table border="1"> <thead> <tr> <th>Code</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00b</td> <td>Indicates the device server supports only Single Transition (ST)</td> </tr> <tr> <td>01b</td> <td>Indicates the device server supports only Double Transition (DT)</td> </tr> <tr> <td>10b</td> <td>Reserved</td> </tr> <tr> <td>11b</td> <td>Indicates the device server supports ST and DT</td> </tr> </tbody> </table>	Code	Description	00b	Indicates the device server supports only Single Transition (ST)	01b	Indicates the device server supports only Double Transition (DT)	10b	Reserved	11b	Indicates the device server supports ST and DT
Code	Description											
00b	Indicates the device server supports only Single Transition (ST)											
01b	Indicates the device server supports only Double Transition (DT)											
10b	Reserved											
11b	Indicates the device server supports ST and DT											
QAS	1 (U160)	Quick Arbitrate Support. If the value of this field is 1, it indicates that the device server supports the quick arbitrate feature. A value of 0 indicates that the device server does not support this feature. 1 is default setting.										
IUS	1 (U320)	Information Unit Supported. If the value of this field is 1, it indicates that the device server supports information units. A value of 0 indicates that the device server does not support information units. Default is 1.										

Note: Vendor Information, Product Identification, and Product Revision Level are returned as shown in Table 5-22.

5.9.2 Vital Product Data Pages

The Vital Product Data pages that can be returned by the drive are described in the following paragraphs in the sequence shown in Table 5-27.

5.9.2.1 Supported Vital Product Data Pages Page (00h)

The Supported Vital Product Data Pages page (Table 5-26) provides a directory of the Vital Product Data Pages that are supported by the drive. Table 5-27 lists the supported pages.

Table 5-26 *Supported Vital Product Data Pages -Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0	Peripheral Qualifier = 0			Peripheral Device Type = 0 (Direct Access Device)				
1	Page Code (00h)							
2	Reserved							
3	Page Length = 08h							
4 - 10	Supported Page List							

Table 5-27 *Vital Product Data-Page Codes*

Page Code	Description	Size In Bytes
00h	Supported Vital Product Pages Page	11
80h	Unit Serial Number Page	12
81h	Implemented Operating Definition Page	11
82h	ASCII Implemented Operating Page	32
83h	Device Identification Page	40

5.9.2.2 Unit Serial Number Page (80h)

The Unit Serial Number page contains the drive's PCB Serial Number (Table 5-28) and the HDA Serial Number (Table 5-29).

Table 5-28 *Unit Serial Number Page-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0	Peripheral Qualifier = 000b			Peripheral Device Type = 000h (Direct Access Device)				
1	Page Code (80h)							
2	Reserved							
3	Page Length = 08							
4 - 11	HDA Serial Number							

Table 5-29 *Unit Serial Number Page-Field Descriptions*

Field	Definition
HDA Serial Number	An 8-character ASCII representation of the drive's HDA serial number

5.9.2.3 Implemented Operating Definition Page (81h)

The Implemented Operating Definition page reflects the current and available operation definitions as described in Table 5-30 and Table 5-31.

Table 5-30 *Implemented Operating Definition Page-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0	Peripheral Qualifier = 0			Peripheral Device Type = 0 (Direct Access Device)				
1	Page Code (81h)							
2	Reserved							
3	Page Length = 07h							
4	Rsvd	Current Definition						
5	SavImp = 1	Default Definition is SCSI-2 U160 (04h)						
6	SavImp = 1	Supported-1 Definition is SCSI-3 U320 (06h)						
7 *	SavImp = 1	Supported-2 Definition is SCSI-3 U160 (05h)						
8	SavImp = 1	Supported-3 Definition is SCSI-2 U160 (04h)						
9	SavImp = 1	Supported-4 Definition is SCSI-2 U80 (03h)						
10	SavImp = 1	Supported-5 Definition is SCSI-1-CCS (02h)						

Table 5-31 *Implemented Operating Definition Page-Field Descriptions*

Field	Definition
SavImp	Save Implemented. If = 0, this bit indicates that the corresponding operation definition cannot be saved. When = 1, this bet indicates that the corresponding operating definition can be saved.

5.9.2.4 ASCII Implemented Operating Definition Page (82h)

The ASCII Implemented Operating Definition page returns the character string's length (1Bh) in byte 4, followed by the appropriate character string (“SCSI-3, SCSI-2, SCSI-1/CCS” in bytes 5 through 31 for Parallel SCSI. This is described in Table 5-32.

Table 5-32 *ASCII Implemented Operating Definition Page - Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0	Peripheral Qualifier = 0			Peripheral Device Type = 0 (Direct Access Device)				
1	Page Code (82h)							
2	Reserved							
3	Page Length = 1Ch							
4	ASCII String Length = 1Bh							
5 - 31	“SCSI-3, SCSI-2, SCSI-1/CCS” + byte of 0							

5.9.2.5 Device Identification Page (83h) (SCSI-3 ONLY)

The Device Identification Page provides the means to retrieve zero or more identification descriptors that apply to the logical unit (Table 5-33 and Table 5-34).

Table 5-33 *Device Identification Page - Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0	Peripheral Qualifier (0)			Peripheral Device Type (0) (Direct Access Device)				
1	Page Code (83h)							
2	Reserved							
3	Page Length (24h for SCSI)							
4	Reserved				Code Set (1)			
5	Reserved		Association (0)		Identifier Type (2)			
6	Reserved							
7	Identifier Length (8h)							
8 – 15	EUI-64 (Value Stored in Configuration Page 30h)							
16	Reserved				Code Set (2)			
17	Reserved		Association (0)		Identifier Type (1)			
18	Reserved							
19	Identifier Length (14h)							
20 – 27	Vendor Identifier (“Maxtor “)							
28 – 39	Drive Serial Number							

Table 5-34 *Device Identification Page - Field Description*

Field	Definition
Code Set	This field specifies the code set used for the Identifier field. Applicable values are: Value Description 0h Reserved 1h The Identifier field contains binary values 2h The Identifier field contains ASCII graphic codes (code values 20h through 7Eh) 3h – Fh Reserved
Association	This field specifies the entity with which the Identifier field is associated. Applicable values are: Value Description 0h The Identifier field is associated with the addressed physical or logical device. 1h The identifier field is associated with the port that received the request. 2h – 3h Reserved
Identifier Type	This field specifies the format and assignment authority for the identifier. Values in this field are: Value Description 0h No assignment authority was used; there is no guarantee that the identifier is globally unique (vendor-specific). 1h The first 8 bytes of the Identifier field represent the Vendor ID. 2h The Identifier field contains an IEEE Extended Unique Identifier, 64-bit (EUI-64). The Identifier Length field (Byte 7) is set to 8. 3h Not applicable; for Fibre Channel devices. 4h Not applicable; if the Association value = 1h, the value of the Identifier contains a 4-byte, binary number that identifies the port relative to the other ports in the device. 5h - Fh Reserved

5.9.2.6 Command Support Data Pages

An application client can request command support data by setting the CmdDt bit of the INQUIRY command to 1, and specifying the SCSI operation code of the Command Descriptor Block (CDB) for which it wants information.

Format of the command support data and definitions of the fields follow in Table 5-35, Table 5-36, and Table 5-37.

Table 5-35 *Command Support Data Page-Data Format*

Bit Byte	7	6	5	4	3	2	1	0	
0	Peripheral Qualifier = 0			Peripheral Device Type = 0 (Direct Access Device)					
1	Reserved					Support			
2	ISO Version		ECMA Version			ANSI-Approved Version			
3 – 4	Reserved								
5	CDB Size (m – 5)								
6 – m	(MSB)			CDB Usage Data					(LSB)

Table 5-36 *Command Support Data Page-Field Descriptions*

Field	Description	
Support	The value of the Support field describes the type of support that the disk drive provides for Command Support Data.	
	<u>Value</u>	<u>Description</u>
	000b	Data about the requested SCSI operation code is not currently available. In this case, all data after Byte 1 is undefined.
	001b	The device does not support the SCSI operation code requested. In this case, all data after Byte 1 is undefined
	010b	Reserved
	011b	The device supports the SCSI operation code in conformance with the SCSI standard.
	100b	Vendor-Specific
	101b	The device supports the SCSI operation code, but in a vendor-specific manner
	110b	Vendor-Specific
	111b	Reserved
ISO-Version	Must be 0.	
ECMA-Version	Must be 0.	
ANSI-Approved Version	Minimum operating definition for supported command.	
CDB Size	This field contains the number of bytes in the CDB for the Operation Code being requested and the size of the CDB Usage Data in the data that is returned in response to the INQUIRY.	
CDB Usage Data	This field contains information about the CDB for the Operation Code being queried. Note that the first byte of the CDB Usage Data contains the OpCode for the operation specified. All of the other bytes of the CDB Usage Data contain a map for bits in the CDB of the OpCode specified.	

Note: The bits in the map have a 1-to-1 correspondence to the CDB for the OpCode being queried. That is, if the device senses a bit as the entire field or as part of the field of the operation, the map in CDB Usage Data contains a 1 in the corresponding bit position. If the device ignores a bit or declares a bit as “reserved” in the CDB for the OpCode being queried, the map has a 0 in that corresponding bit position. Refer to Table 5-36 for a list of the data returned for each of the OpCode values that can be sent in the INQUIRY command.

Table 5-37 *Command Support Data Page Command or Operation Codes*

		Hex Data Returned When INQUIRY is Received and CmdDt Bit = 1			
OpCode	Command	Support	ANSI Version	CDB Length	CDB Size Usage Data
00h	TEST UNIT READY	03	01	06	00E000000003
01h	REZERO UNIT	03	01	06	01E000000003
03h	REQUEST SENSE	03	01	06	03E00000FF03
04h	FORMAT UNIT	03	01	06	04FFFFFFFFF03
07h	REASSIGN BLOCKS	03	01	06	070E00000003
08h	READ (6)	03	01	06	08FFFFFFFFF03
0Ah	WRITE (6)	03	01	06	0AFFFFFFFFF03
0Bh	SEEK (6)	03	01	06	0BFFFFFFFF0003
12h	INQUIRY	03	01	06	12E3FF00FF03
15h	MODE SELECT (6)	03	01	06	15F10000FF03
16h	RESERVE (6)	03	01	06	16FEFFFFFFFF03
17h	RELEASE (6)	03	01	06	17FEFF000003
1Ah	MODE SENSE (6)	03	01	06	1AE8FF00FF03
1Bh	START STOP UNIT	03	01	06	1BE100000103
1Ch	REC. DIAG. RESULTS	03	01	06	1CE1FFFFFFFF03
1Dh	END DIAGNOSTIC	03	01	06	1DF700FFFF03
25h	READ CAPACITY	03	01	0A	25E0FFFFFFFF00000103
28h	READ (10)	03	01	0A	28F8FFFFFFFF00FF03
2Ah	WRITE (10)	03	01	0A	2AF8FFFFFFFF00FF03
2Bh	SEEK (10)	03	01	0A	2BE0FFFFFFFF00000003
2Eh	WRITE AND VERIFY	03	01	0A	2EF2FFFFFFFF00FF03
2Fh	VERIFY	03	01	0A	2FF2FFFFFFFF00FF03
35h	SYNCH. CACHE	03	01	0A	35E0FFFFFFFF00FF03
37h	READ DEFECT DATA	03	01	0A	37E01F00000000FF03
3Bh	WRITE BUFFER	03	01	0A	3BEFFFFFFFFE00FF03
3Ch	READ BUFFER	03	01	0A	3CEFFFFFFFFFFFF03
3Eh	READ LONG	03	01	0A	3EE2FFFFFFFF00FF03
3Fh	WRITE LONG	03	01	0A	3FE0FFFFFFFF00FF03

		Hex Data Returned When INQUIRY is Received and CmdDt Bit = 1			
40h	CHANGE DEF.	03	01	0A	40E0017F000000000003
41h	WRITE SAME	03	01	0A	41E2FFFFFFFFF00FF FF03
4Ch	LOG SELECT	03	01	0A	4CE3C000000000FF FF03
4Dh	LOG SENSE	03	01	0A	4DE1FF00000000FF FF03
55h	MODE SELECT (10)	03	01	0A	55100000000000FFF F03
56h	RESERVE (10)	03	03	0A	5610FFFF000000FF FF03
57h	RELEASE (10)	03	03	0A	5710FFFF000000FF FF03
5Ah	MODE SENSE (10)	03	01	0A	5AE8FF00000000FF FF03
5Eh	PERSIST. RES. IN	03	03	0A	5E1F0000000000FF FF03
5Fh	PERSIST. RES. OUT	03	03	0A	5F1FFF00000000FF FF03
A0h	REPORT LUNS	03	03	0C	A00000000000FFFF FFFF0003

5.10 LOG SELECT Command (4Ch)

The drive collects and stores performance data and error summaries in counters. The LOG SELECT command is used to zero these counters. The LOG SELECT command is a complementary command to the LOG SENSE command. The format of the LOG SELECT CDB and a description of the fields follows in Table 5-38 and Table 5-39 respectively.

Table 5-38 LOG SELECT Command Descriptor Block-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (4Ch)							
1	Reserved						PCR	SP
2	PC		Reserved					
3 – 6	Reserved							
7 – 8	Parameter List Length							
9	Control							

Table 5-39 LOG SELECT Command Descriptor Block-Field Descriptions

Field	Description
PCR	Parameter Code Reset. Must be 1. Causes all the implemented counters to be set to 0.
SP	Save Parameters. A value of 1 indicates that certain counters are to be cleared from non-volatile memory.
PC	Page Control. Must be 01b OR 11b.
Parameter List Length	Must be 0.

5.11 LOG SENSE Command (4Dh)

Note: Log Sense data pages require special interpretation and also are subject to change. For assistance with the Log Sense data pages, contact your Maxtor Applications Engineer.

The drive collects operational information and stores these statistics as *log data*. Log data are grouped by category into *log pages*. The LOG SENSE command allows an initiator to retrieve the stored log data. The LOG SENSE command is a complementary command to the LOG SELECT command.

Table Table 5-40 lists the log pages supported by the drive. Contact your Maxtor Applications Engineer for more information.

Table 5-40 *Disk Drive Log Pages*

Page Code	Description
00h	Supported Log Pages
01h	Buffer Overruns and Underruns
02h	Write Error Counter
03h	Read Error Counter
05h	Verify Error Counter
06h	Non-Medium Error Counter Page
07h	Last n-Error Events Page
08h	Format Status Page
0Dh	Temperature Page
0Eh	Start-Stop Cycle Counter
0Fh	Application Client Page
10h	Self Test Results Page
2Fh	EWS Status

5.11.1 LOG SENSE Command Descriptor Block

The Command Descriptor Block for the LOG SENSE command is shown in Table 5-41. Table 5-42 contains field descriptions.

Table 5-41 LOG SENSE Command Descriptor Block-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (4Dh)							
1	Reserved						PPC	SP
2	PC		Page Code					
3 – 4	Reserved							
5 – 6	Parameter Pointer							
7 – 8	Allocation Length							
9	Control							

Table 5-42 LOG SENSE Command Descriptor Block-Field Descriptions

Field	Description
PPC	The Parameter Pointer Control bit controls the type of parameters that can be requested from the drive. This bit must be 0, indicating that all log parameters for the specified page (subject to the allocation length specified) are returned to the initiator.
SP	The Save Pages bit specifies whether the parameters are to be saved. This bit must be 0, indicating that no parameters are to be saved and are reset at power-on or by a TARGET RESET.
PC	The Page Control field defines the type of parameter values to be selected. The field must be 01b indicating that the current values are to be returned. Mode 11b (return default values) is not supported because all counters have a default value of 0.
Page Code	The value specified as the Page Code determines the page to be returned. Table 5-40 contains a list of supported log pages and their page codes.
Parameter Pointer	This field is related to the PPC field. This bit must be 0.
Allocation Length	This specifies the number of bytes of data that the drive is allowed to pass during the DATA IN phase. The requested page is truncated if its length exceeds the number of bytes specified in this field.

5.11.2 LOG SENSE Log Pages

The log pages that are returned from the drive have a common format that is shown in Table 5-43. Each page contains a 4-byte header followed by one or more log parameters. Refer to Table 5-44. Table 5-45 contains the format for the Generic Log Parameter and Table 5-46 describes the fields.

Table 5-43 LOG SENSE Log Page Format-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0	Reserved		Page Code					
1	Reserved							
2 – 3	Page Length (n-3)							
4 to x+3	Log Parameter (first) (length x) ¹							
n-y+1 to n	Log Parameter (last) (length y) ¹							

Note: ¹Length x or y is the sum of parameter header and parameter data value bytes.

Table 5-44 LOG SENSE Log Page Format-Field Descriptions

Field	Description
Page Code	The page code as given in Table 5-39.
Page Length	The allocation length for the page minus the 4-byte header.
Log Parameter	One or more log data entities that are returned as part of a Log Page. Each Log Parameter starts with a 4-byte header followed by one or more bytes of <i>value</i> data. In most cases, the parameter value is a 4-byte longword that contains the present value of an error or performance counter. See Table 5-45 for the format of a generic log parameter.

Table 5-45 Generic Log Parameter-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0 – 1	Parameter Code							
2	DU	DS	TSD	ETC		TMC	LBIN	LP
3	Page Length (n-3)							
4 – n	Parameter Value							

Table 5-46 *Generic Log Parameter-Field Descriptions*

Field	Description
Parameter Code	A code which uniquely identifies each parameter on a given Log Page. For example, the code 8002h on the Seek Performance Summary page reports the average seek time while the same code on the Block Replacement Summary page reports the number of blocks replaced. The following pages in this manual contain the lists of applicable parameters for each Log page.
DU	Disable Update. This parameter is 1 when updates are not enabled.
DS	Disable Save. This parameter is 1 when parameters are not saved by the drive.
TSD	Target Save Disable. This parameter is 0 indicating that the drive provides a target-specific way of saving parameters.
ETC	Enable Threshold Comparison. This parameter is always 0. The drive does not use thresholds for any of its parameters.
TMC	Threshold Met Comparison. This parameter is always 0. The drive does not use thresholds for any of its parameters.
LBIN	List in Binary. This bit only valid if LP = 1. If LP = 1 and LBIN = 0, then the List Parameter is a string of ASCII graphic codes (code values 20h through 7Eh). If LP = 1 and LBIN = 1, then the parameter is a list of binary information.
LP	List Parameter. When 0, indicates that the parameter value for this parameter is a numeric <i>value</i> . When LP is a 1, it indicates that the parameter value is an alphanumeric ASCII-string <i>list</i> . This parameter is always 0.
Parameter Length	Specifies the length, in bytes, of the parameter's value.
Parameter Value	Contains the parameter's current value when the PC field of the LOG SENSE command is 01b.

Note: The DU, DS, TSD, ETC, TMC, LBIN and LP fields are collectively referred to as the Parameter Control Byte. This byte generally has a value of 0; however, if the page is non-volatile but the drive is not spun up or the GLTSD bit is set in Mode Page 0Ah, then the value would be 20h. If the page is volatile, then the value is 60h.

5.12 MODE SELECT (6) Command (15h)

SCSI refers to the drive's operational parameters as its *mode parameters*. SCSI groups the mode parameters by function into a set of data structures referred to as *mode pages*. The MODE SELECT (6) command allows the initiator to modify some of these mode pages and thereby control some of the drive's operational characteristics. The Save Page (SP) option in the Command Descriptor Block makes the changes permanent. The new mode parameters are then stored in the drive's non-volatile memory. The MODE SELECT CDB is described in Table 5-47 and a description of the fields in Table 5-48.

Table 5-47 *MODE SELECT (6) Command Descriptor Block-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (15h)							
1	Reserved			PF	Reserved			SP
2 – 3	Reserved							
4	Parameter Length List							
5	Control							

Table 5-48 *MODE SELECT (6) Command Field Descriptions*

Data Field	Description
PF	Page Format. The drive ignores the content of this field and produces a mode parameter list that contains a mode header optionally followed by a sector descriptor and one or more pages that conform to the format shown here.
SP	Save Pages. When SP = 0, the drive performs the specified MODE SELECT operation and does not save any pages. When SP = 1, the drive performs the specified operation and stores all savable pages, (PS=1 on MODE SENSE return) including any sent during the data-out buffer transfer. The changes made when SP = 1 become permanent changes to the drive's SCSI setup.
Parameter List Length	This field tells the drive how many bytes of Mode Parameters to fetch in the data-out buffer transfer. A CHECK CONDITION status is returned if this value truncates a page.

5.12.1 Initiator-Changeable Mode Pages

Table 5-49 lists the mode pages that are supported by the drive. An initiator can change these pages by supplying them, with the desired changes included, in the MODE SELECT command's data-out buffer transfer.

The initiator should first use a MODE SENSE command to read the appropriate pages and leave non-changeable values as read when the initiator subsequently writes the changeable pages with the MODE SELECT command. If a non-changeable field contains an invalid value, the drive returns a CHECK CONDITION status.

Table 5-49 *Initiator-Changeable Mode Pages*

Page Code	Page Name	Function	Size In Bytes
00h	Unit Attention Control Page	Unit Attention reporting (enables or disables)	4
01h	Read-Write Error Recovery Page	Medium Access Error recovery and reporting procedures for READ and WRITE commands	12
02h	Disconnect-Reconnect Page	Bus behavior during data transfers	16
07h ¹	Verify Error Recovery Page	Medium Access Error recovery and reporting procedures for the VERIFY command	12
08h ¹	Caching Page	Cache policy	20
0Ah ¹	Control Mode Page	Command processing policy	12
0Ch	Notch and Partition Page	Drive geometry reporting	24
19h	Port Control Page	Defines port control parameters	8
19h(1h)	Margin Control	Margin control values	16
19h(2h)	Saved Training Configuration	Saved training configuration values	234
19h(3h)	Negotiated Settings	Negotiated settings for current I_T Nexus	14
19h(4h)	Report Transfer Capabilities	Transfer capabilities	14
1Ah	Power Condition Page	Enable and set time intervals for Idle and Standby modes	12
1Ch ¹	Information Exceptions Control Page and operations of specific informational exception conditions.	Defines methods to control reporting	12

NOTE:

1. This page is not supported by the SCSI-1/CCS Operating Mode.

5.12.2 Mode Page Types

The drive maintains three distinct sets of mode pages. They are the *current page*, the *default page*, and the *saved page*. The drive also reports a fourth set of *changeable* pages. The page types are defined in Table 5-50.

Table 5-50 *Mode Page Types*

Page Type	Definition
Current	<p>The <i>current</i> mode page set applies to all <i>initiators</i> and defines the drive's mode. The SCSI-2 specification states that a drive can maintain Mode parameters on a per-initiator basis if it so chooses. The Drives do not support this option.</p> <p>The <i>current</i> mode page set contains the values supplied in the last MODE SELECT command received from an initiator. If no initiator has sent a MODE SELECT command since the drive was last reset or powered up, the current pages contain</p> <ul style="list-style-type: none"> • <i>Saved</i> values if saved pages exist (from a previous MODE SELECT command SP parameter). • <i>Default</i> values if pages have never been saved. <p>The drives generate a Unit Attention condition for all initiators (except for the one that was the source of the MODE SELECT command) whenever one initiator modifies the Mode parameters.</p>
Default	The <i>default</i> mode page set contains the factory default values that are listed in each page's description.
Saved	The <i>saved</i> mode page set contains values preserved in the drive's non-volatile memory by a previous SP-modified MODE SELECT command.
Changeable	The <i>changeable</i> mode page set provides a means for an initiator to determine which pages it is allowed to change and the specific bits within those pages that it is allowed to change. This page set is read-only and is fetched with a MODE SENSE command.

5.12.3 Mode Parameter List

Table 5-51 shows the format of the Mode Parameter List that is passed by the initiator to the drive during the command's data-out buffer transfer. Table 5-52 describes the fields. Table 5-53 and Table 5-54 provides a description of the data format and fields of the Mode Parameter Header. Table 5-55 and Table 5-56 describe the format and field descriptions for the Mode Parameter Block Descriptor.

Table 5-51 *Mode Parameter List-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0 – 3	Mode Parameter Header							
4 – 11	Block Descriptor (optional)							
4 – n or 12 – n	Page(s) (optional)							

Table 5-52 *Mode Parameter List-Field Descriptions*

Field	Description
Mode Parameter Header	Contains information about the remainder of the parameter list and is always present (see Table 5-53 and Table 5-54).
Block Descriptor	Allows the initiator to set the drive's Logical Block Size and number of Logical Block Addresses (see Table 5-55 and Table 5-56).
Page(s)	The page code(s) of the pages that are a part of this command.

Table 5-53 *Mode Parameter Header (6-Byte)-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0	Mode Data Length							
1	Medium Type							
2	Device-Specific Parameter							
3	Block Descriptor Length							

Table 5-54 *Mode Parameter Header- Field Descriptions*

Field	Description
Mode Data Length	Reserved. Must be 0.
Medium Type	Ignored by the drive.
Device-specific Parameter	Ignored by the drive.
Block Descriptor Length	Zero (0) if no Block Descriptor is supplied. The length is eight (8) if a Block Descriptor is supplied.

Table 5-55 *Mode Parameter Block Descriptor-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0 – 3	Number of Blocks							
4	Reserved							
5 – 7	Block Length							

Table 5-56 *Mode Parameter Block Descriptor-Field Descriptions*

Field	Description
Number of Blocks	<p>If the number of blocks is set to zero, the device shall retain its current capacity if the block size has not changed. If the number of blocks is set to zero and the block size has changed, the device shall be set to its maximum capacity when the new block size takes effect.</p> <p>If the number of blocks is greater than zero and less than or equal to its maximum capacity, the device shall be set to that number of blocks. If the block size has not changed, the device shall not become format corrupted. This capacity setting shall be retained through power cycles, hard resets, logical unit resets and I_T nexus losses:</p> <p>If the number of blocks field is set to a value greater than the maximum capacity of the device and less than FFFF FFFFh, then the command is terminated with a CHECK CONDITION status. The sense key is set to ILLEGAL REQUEST. The device shall retain its previous block descriptor settings; or</p> <p>If the number of blocks is set to FFFF FFFFh, the device shall be set to its maximum capacity. If the block size has not changed, the device shall not become format corrupted. This capacity setting shall be retained through power cycles, hard resets, logical unit resets, and I_T nexus losses.</p>
Block Length	This field specifies the length, in bytes, of each logical sector.

5.12.4 Categories of Changeable Pages

The drive's changeable pages are described on the following pages. The data fields for each of these pages fall into one of the categories described in Table 5-57.

Table 5-57 *Categories of Changeable Pages*

Category	Description
Fully Supported	A field in this category can be modified by the MODE SELECT command, can be saved and can affect the drives processing. The drive uses the value in the field and, if the SP bit is set, preserves the new value of the field in non-volatile memory on the drive. (The PS bit must have been set when the mode page was returned with a MODE SENSE command.)
Ignored	A value in this field is never used or validated; it is never looked at by the drive. Ignored fields are not underlined in the Mode Parameter Pages' figures or tables and are not described in the Field Description tables.

5.12.5 Unit Attention Condition Page (00h)

The Unit Attention Page is used to enable or disable the generation of Unit Attention Conditions. The Unit Attention Condition Page format and field descriptions are described in Table 5-58 and Table 5-59 respectively.

Table 5-58 *Unit Attention Condition Page (Page 0)*

Bit Byte	7	6	5	4	3	2	1	0
0	PS	Rsvd	Page Code (00h)					
1	Page Length (02h)							
2	Reserved			DUA	Reserved			
3	Reserved (00h)							

Table 5-59 *Unit Attention Condition Page (Page 0)*

Field	Default Value	Description
PS	1	Parameters Savable. This bit is only used with the MODE SENSE command. A returned value of 1 indicates that the target is capable of saving the page in a non-volatile, vendor-specific location. The bit is reserved with the MODE SELECT command.
DUA	0	Disable Unit Attention. When set to the default value of 0, the drive reports a CHECK CONDITION for any UNIT ATTENTION condition (power-on, SCSI bus reset, etc.)

5.12.6 Read-Write Error Recovery Page (01h)

The Read-Write Error Recovery page controls the drive's response to error conditions that arise during the course of READ (6), READ (10), WRITE (6), WRITE (10) command processing and during the write portion of WRITE AND VERIFY command processing. The Read-Write Error Recovery page format and field descriptions are described in Table 5-60 and Table 5-61 respectively.

Table 5-60 *Read-Write Error Recovery Page-(Page 1)*

Bit Byte	7	6	5	4	3	2	1	0
0	PS	Rsvd	Page Code (01h)					
1	Page Length (0Ah)							
2	AWRE	ARRE	TB	RC	EER	PER	DTE	DCR
3	Read Retry Count							
4	Correction Span							
5	Head Offset Count							
6	Data Strobe Offset Count							
7	Reserved							
8	Write Retry Count							
9	Reserved							
10 – 11	Recovery Time Limit							

Table 5-61 *Read-Write Error Recovery Page - Field Descriptions*

Field	Default Value	Description
PS	1	Parameters Savable. This bit is only used with the MODE SENSE command. A returned value of 1 indicates that the target is capable of saving the page in a non-volatile, vendor-specific location. The bit is reserved with the MODE SELECT command.
AWRE	1	Automatic Write Reallocation Enable. When the value is 1, it causes the drive to automatically reallocate bad sectors encountered during a write operation. If it succeeds in replacing the bad sector, the drive adds the revectorized sector's Logical Block Number to the Grown Defect List. When AWRE is 0, the drive does not automatically reallocate bad sectors encountered during a write operation.
ARRE	1	Automatic Read Reallocation Enable. When the value is 1, it causes the drive to automatically reallocate bad sectors encountered during a read operation. If it succeeds in replacing the bad sector, the drive adds the revectorized sector's Logical Block Number to the Grown Defect List. Reallocation is done only when the drive is able to recover the sector's data unless the Reallocate Uncorrected Errors bit of the Maxtor (Vendor) Special Function Control Page is set. When ARRE is 0, the drive does not automatically reallocate bad sectors encountered during a read operation. See also Maxtor-Unique Page 39h (RUEE).
TB	0	Transfer Block. When the value is 1, the drive sends best-guess information to the initiator when it encounters a defective sector whose data cannot be fully recovered. When the value is 0, data that cannot be fully recovered is not returned to the initiator.
RC	0	Read Continuous. When the value is 0, the drive attempts error recovery when it encounters an error. When the value is 1, the drive suppresses error recovery. <i>The drive ignores any value placed in this field.</i>
EER	0	Enable Early Recovery. <i>Not supported by the drive.</i>
PER	0	Post Error Recovery. When the value is 0, the drive returns a GOOD status in cases where error recovery is invoked and successful. When the value is 1, the drive returns a CHECK CONDITION status for recovered errors, sets the sense key to RECOVERED ERROR, and posts the appropriate additional sense key.
DTE	0	Disable Transfer on Error. <i>Not supported by the drive.</i>
DCR	0	Disable Correction. When the value is 0, ECC correction is applied as a means of recovering erred data. On write operations, a valid ECC is written to each data sector regardless of the value in the DCR bit. When the value is 1, Firmware ECC correction only is disabled (but not retries or reads), and ECC errors greater than hardware correctable cannot be corrected.
Read Retry Count	4	Specifies the number of retries attempted for a failed read operation.
Correction Span	240	Specifies the maximum number of bits in a data sector that the drive is allowed to correct. A value less than the default value can cause an otherwise correctable error to be reported as uncorrectable. If the initiator specifies a value larger than the default value, the drive automatically rounds down the number to the default value. A value of 0 causes the drive to use the default value. The drive ignores any value placed in this field.
Head Offset Count	0	Not supported by the drive.
Data Strobe Offset Count	0	Not supported by the drive.
Write Retry Count	8	Specifies the number of retries attempted for a failed write operation.

Field	Default Value	Description
Recovery Time Limit	0	Specifies the maximum time, in milliseconds, that a retry is attempted on a failed sector during a read or write operation. When the value is 0, it means that there is no time limit. The minimum permissible value, however, is 100 milliseconds.

5.12.7 Disconnect–Reconnect Page (02h)

The Disconnect–Reconnect Page provides the application client the means to attempt to optimize the performance of the delivery subsystem. The data format and field descriptions are described in Table 5-62 and Table 5-63.

Table 5-62 *Disconnect–Reconnect Page (Page 2)*

Byte	Bit	7	6	5	4	3	2	1	0
0		PS	Rsvd	Page Code (02h)					
1		Page Length (0Eh)							
2		Buffer Full Ratio							
3		Buffer Empty Ratio							
4 – 5		Bus Inactivity Limit							
6 – 7		Disconnect Time Limit							
8 – 9		Connect Time Limit							
10 – 11		Maximum Burst Size							
12		EMDP	Fairness Arbitration			DImm	DTDC		
13		Reserved							
14 – 15		First Burst Size							

Table 5-63 *Disconnect-Reconnect-Field Descriptions (Page 2)*

Field	Default Value	Description
PS	0	Parameters Savable. This bit is only used with the MODE SENSE command. A returned value of 1 indicates that the target is capable of saving the page in a non-volatile, vendor-specific location. The bit is reserved with the MODE SELECT command.
Buffer Full Ratio	0	The Buffer Full Ratio indicates how full the buffer should be prior to requesting an interconnect tenancy. Applies to READ operations only. The default of 0 indicates to the drive that it should use its default ratio.
Buffer Empty Ratio	0	The Buffer Empty Ratio indicates how full the buffer should be prior to requesting an interconnect tenancy. Applies to READ operations only. The default of 0 indicates to the drive that it should use its default ratio. <i>Not supported by the drive.</i>
Bus Inactivity Limit	0	The Bus Inactivity Limit indicates the maximum duration of any interconnect tenancy during which no data is transferred. <i>Not supported by the drive.</i>
Disconnect Time Limit	0	The Disconnect Time Limit specifies the minimum time, in 100 microsecond increments, that the drive waits after releasing the SCSI bus before reselecting. The drive supports a maximum value of 0xFF. Values greater than 0xFF are treated as 0xFF. The default value of 0 indicates that the drive can reslect immediately after releasing the bus.
Connect Time Limit	0	The Connect Time Limit indicates the maximum duration of a single interconnect tenancy. <i>Not supported by the drive.</i>
Maximum Burst Size	0	The Maximum Burst Size specifies the maximum amount of data that the drive will transfer during a data phase before disconnecting from the bus, assuming it was given the disconnect privilege. This value is expressed in increments of data sectors (that is, a value of 1 indicates 512 bytes, a value of 2 indicates 1024 bytes). The default value of 0 indicates that there is no limit on the amount of data transferred per connection. <i>This field is reserved for the SCSI-1/CCS operating mode.</i>

Field	Default Value	Description
EMDP	0	Enable MODIFY DATA POINTER. <i>Not supported by the drive.</i>
Fairness Arbitration	111b (Def)	Indicates whether or not the target should use fair or unfair (e.g., priority) arbitration when beginning the interconnect tenancy.
		<u>000b</u> - Disables the fairness algorithm. A fixed priority scheme based on the SCSI ID assigned is used.
		<u>Non-Zero</u> - Any non-zero value enables the fairness algorithm (SCSI SPI-2 ANSI standard, Appendix B).
DImm	0	Disconnect Immediate. When set to a 1, to reduce on-bus time, the drive is forced to disconnect from the SCSI Bus after a SCSI Command is received. Disconnect Immediate works only when a typical, error-free command is received. Disconnect Immediate does not apply in exception conditions (for example, SDTR received, error, etc.).
DTDC	Data Transfer Disconnect Control. This field of three bits defines further restrictions when a disconnect is permitted. These are as follows:	
	Value	Description
	000b (Def)	Data transfer disconnect control is not used. Disconnect is controlled by other fields in this page.
		<u>001b</u> - A target will not attempt to disconnect once the data transfer of a command has started until all data of the command has been transferred. The connect time limit and bus inactivity limit are ignored during the transfer.
		<u>010b</u> - Reserved
		<u>011b</u> - A target will not attempt to disconnect once the data transfer of a command has started, until the command is complete. The connect time limit and bus inactivity limit are ignored once the transfer has started.
First Burst Size	0	This field indicates the maximum amount of data that a target may transfer for a command during the same interconnect in which it receives the command. Not applicable for parallel SCSI.

5.12.8 Verify Error Recovery Page (07h)

The Verify Error Recovery Page controls the drive's response to error conditions that arise during the VERIFY command and during the verify operation of the WRITE AND VERIFY command. The Verify Error Recovery Page is not supported in the SCSI-1/CCS Operating Mode. The data format and field descriptions are described in Table 5-64 and Table 5-65 respectively.

Table 5-64 *Verify Error Recovery Page-(Page 7)*

Byte	Bit	7	6	5	4	3	2	1	0
0		PS	Rsvd	Page Code (07h)					
1		Page Length (0Ah)							
2		Reserved				EER	PER	DTE	DCR
3		Verify Retry Count							
4		Verify Correction Span							
5 – 9		Reserved							
10 – 11		Verify Recovery Time Limit							

Table 5-65 *Verify Error Recovery Page-Field Descriptions (Page 7)*

Field	Default Value	Description
PS	0	Parameters Savable. This bit is only used with the MODE SENSE command. A returned value of 1 indicates that the target is capable of saving the page in a non-volatile, vendor-specific location. The bit is reserved with the MODE SELECT command.
EER	0	Enable Early Recovery. <i>Not supported by the drive.</i>
PER	0	Post Error Recovery. When the value is 1, the drive returns CHECK CONDITION status for recovered errors, sets the sense key to RECOVERED ERROR, and posts the appropriate additional sense code. When set to the default value of 0, it causes the drive to return GOOD status in cases where error recovery is invoked and is successful. The only recoverable error that can occur during a VERIFY operation is a failed read that succeeds after retries.
DTE	0	Disable Transfer on Error. <i>Not supported by the drive.</i>
DCR	0	Disable Correction. When the value is 1, it disables ECC Firmware only (but not retries) on the READ portion of a VERIFY operation. When set to the default value of 0, the ECC correction is always applied as a means of recovering erred data.
Verify Retry Count	4	The Verify Retry Count specifies the number of retries that are attempted per sector on a failed READ portion of a VERIFY operation.
Verify Correction Span	240	The Verify Correction Span specifies the maximum number of bits in a sector that the drive is allowed to correct. A value less than the default value can cause an otherwise correctable error to be reported as uncorrectable. If the initiator specifies a value larger than 80, the drive rounds down to 80 without reporting the change. A value of 0 causes the drive to use its default value. The drive ignores any value in this field.
Verify Recovery Time Limit	0	The Verify Recovery Time Limit specifies (in milliseconds) the maximum time that a retry is attempted on a failed sector during the read portion of a verify operation. The default value of 0 indicates that there is no time limit. The minimum permissible is 100 milliseconds.

5.12.9 Caching Page (08h)

The Caching Page controls the drive's cache management policy. The data format and field descriptions are described in Table 5-66 and Table 5-67 respectively.

Table 5-66 *Caching Page-(Page 8)*

Byte	Bit	7	6	5	4	3	2	1	0
0		PS	Rsvd	Page Code (08h)					
1		Page Length (12h)							
2		IC	ABPF	CAP	DISC	Size	WCE	MF	RCD
3		Demand Read Retention Priority				Demand Write Retention Priority			
4 – 5		Disable Pre-Fetch Transfer Length							
6 – 7		Minimum Pre-Fetch							
8 – 9		Maximum Pre-Fetch							
10 – 11		Maximum Pre-Fetch Ceiling							
12		FSW	LBCSS	DRA	VS	VS	Reserved		
13		Number of Cache Segments							
14 – 15		Cache Segment Size							
16		Reserved							
17 – 19		Non-Cache Segment Size							

NOTES:

Bytes 12 to 19 are returned only if the device operating definition is SCSI-3.

Table 5-67 *Caching Page - Field Descriptions*

Field	Default Value	Description
PS	1	Parameters Savable. This bit is only used with the MODE SENSE command. A returned value of 1 indicates that the target is capable of saving the page in a non-volatile, vendor-specific location. Bit reserved with MODE SELECT command.
IC	0	An initiator control (ic) enable bit set to one specifies that the device server use the <i>NUMBER OF CACHE SEGMENTS</i> field or the <i>CACHE SEGEMENT SIZE</i> field, dependent upon the <i>SIZE</i> bit, to control the caching algorithm rather than the device server's own adaptive algorithm.
ABPF	0	Enable Abort Prefetch. When set, this bit requests the device abort pre-fetch work on the receipt of a new command. <i>Ignored by the drive</i>
CAP	0	A caching analysis permitted (cap) bit set to one specifies that the device server perform caching analysis during subsequent operations. A cap bit set to zero specifies that the caching analysis be disabled to reduce overhead time or to prevent non pertinent operations from impacting tuning values.
DISC	0	A discontinuity (DISC) bit set to one specifies that the device server continue the pre-fetch across time discontinuities (e.g., across cylinders) up to the limits of the buffer, or segment, space available for the pre-fetch. A DISC bit set to zero specifies that pre-fetches be truncated or wrapped at time discontinuities. <i>ignored by drive.</i>
Size	0	Specifies which of two fields are used to configure the cache. When Size=0, the Number of Cache Segments field is used. When Size=1, the Cache Segment Size field is used.
WCE	1	Write Cache Enable. The default value of 1 enables write-back caching when processing a WRITE command. When set to 0 the drive returns a GOOD status and COMMAND COMPLETE only after successfully writing all the data to the media. When WCE = 1, write-back caching is enabled while processing a WRITE command. The drive returns a GOOD status and COMMAND COMPLETE only after fetching the data from the initiator and placing it in the drive's cache memory. Note: For maximum data integrity, it is <u>strongly recommended</u> that when WCE is set to 1, the Read-Write Error Recovery Page (01h) fields be set as follows: <ul style="list-style-type: none"> • AWRE to 1 • RC to 0 • DCR to 0 • Correction Span set to 170 • Head Offset to 0 • Servo Offset to 0
MF	0	Multiplication Factor (value must be 0). The drive interprets the minimum and maximum pre-fetch fields values in terms of the number of logical blocks for each of the respective types of prefetch.
RCD	0	Read Cache Disable. The default value of 0 enables the read cache capability and allows the drive to use cache-resident data or medium data to satisfy a READ command. When the value is 1, read caching is disabled and the data returned in response to a READ command is fetched from the medium and not from the drive's cache memory.

Field	Default Value	Description
Demand Read Retention Priority	0	This field advises the drive on the retention priority to assign data read into the cache that has also been transferred from the drive to the initiator.
Demand Write Retention Priority	0	This field advises the drive on the retention priority to assign data written into the cache that has also been transferred from the cache memory to the medium.
Disable Pre-Fetch Transfer Length	FFFFh	This field specifies the selective disabling of anticipatory pre-fetch on long transfer lengths. The value in this field is compared to the number of sectors requested by the current READ command. If the number of sectors is greater than the value specified, anticipatory pre-fetch is not performed for the command. When 0, anticipatory pre-fetch is disabled.
Minimum Pre-fetch	0x100	This field specifies the desired number of sectors to be pre-fetched preceding the requested Data.
Maximum Pre-fetch	0x200	This field specifies the desired amount of total number of sectors to be pre-fetched surrounding the requested data.
Maximum Pre-fetch Ceiling	0x200	Specifies an upper limit on the number of logical blocks computed as the maximum pre-fetch. If this number of sectors is greater than the maximum pre-fetch, then the number of logical blocks to pre-fetch is reduced to the value stored in the maximum pre-fetch ceiling field.
FSW	0	The Force Sequential Write bit. <i>Not supported by the drive.</i>
LBCSS	TBD	Selects the units of the cache segment Size field. When LBCSS=0, Units are bytes per segment. When LBCSS=1, Units are blocks per segment.
DRA	0	When set, the device server will not read into cache any logical blocks beyond the addressed logical block (s). When not set, the device server may continue to read logical blocks beyond the addressed logical block (s).
VS		The Vendor-Specific bits. <i>Not supported by the drive.</i>
Number of Cache Segments	variable	Specifies the number of cache segments to divide the cache into. This field is used when the Size bit is 0.
Cache Segment Size	variable	Specifies the cache segment size, in either bytes per segment or blocks per segment, depending on the state of the LBCSS bit. This field is used when the Size bit is 1
Non-Cache Segment Size	0	<i>Not supported by the drive.</i>

5.12.10 Control Mode Page (0Ah)

The Control Mode page controls the drive's command processing and error handling policies. The data format and field descriptions are described in Table 5-68 and Table 5-69 respectively.

Table 5-68 Control Mode Page - Data Format (Page 0Ah)

Bit Byte	7	6	5	4	3	2	1	0
0	PS	Rsvd	Page Code (0Ah)					
1	Page Length (0Ah)							
2	Reserved						GLTSD	RLEC
3	Queue Algorithm Modifier				Re-served	Qerr		DQue

4	EECA	RAC	Reserved	SWP	RAERP	UAAER P	EAERP	
5	Reserved							
6 – 7	Ready AEN Holdoff Period							
8 – 9	Busy Timeout Period							
10	MSB	Extended Selftest Completion Time						
11	Extended Selftest Completion Time						LSB	

NOTES:

Bytes 8 to 11 are returned only if the device operating definition is SCSI-3.

Table 5-69 Control Mode Page-Field Descriptions

Field Value	Default	Description
PS	1	Parameters Savable. This bit is only used with the MODE SENSE command. A returned value of 1 indicates that the target is capable of saving the page in a non-volatile, vendor-specific location. The bit is reserved with the MODE SELECT command.
GLTSD	0	Global Logging Target Save Disable. If this bit = 1, the ability to save log ages to disk is disabled. If a LOG SELECT with the SP bit is received by the drive, the current cumulative values will be cleared. The on-disk values will not be changed. If a LOG SENSE with the SP bit set is received by the drive and the GLTSD bit is set and the unit is ready, the last LOG counter values written to disk will be returned (not the current cumulative values). If a STOP UNIT command or a LOG SENSE command with the SP bit set is issued to the drive, the cumulative values of the Log Counters page will be written to disk. A LOG SENSE command will return the cumulative values since the counters were cleared, not since the drive was powered on.
RLEC	0	Report Log Exception Condition. The drive does not report log exception conditions. <i>Not supported by the drive.</i>
Queue Algorithm Modifier	0001b	This field specifies restrictions on the algorithm used for re-ordering commands that are tagged with the SIMPLE message. A value of 0 specifies that the drive orders the actual execution sequence with a SIMPLE statement. A value of 1 (Unrestricted Re-ordering Allowed) specifies that the drive can re-order the actual execution sequence of the commands with a SIMPLE in any manner. Any data integrity exposures related to command sequence order are explicitly handled by the initiator through the selection of appropriate commands and queue tag messages.
QErr	00b	Queue Error Management. This field controls the drive's processing of queued commands when a Contingent Allegiance (CA) condition or Extended Contingent Allegiance (ECA) is cleared. Processing of queued commands is resumed when a CA/ECA condition is cleared.
DQue	0	Disable Queue. When set to the default value of 0, the drive supports tagged command queuing. When the Disable Queue bit is set to 1, tagged command queuing is disabled. A subsequent tagged message from the initiator is rejected with a MESSAGE REJECT message and the I/O process is executed as an untagged command.
EECA	0	AEN enabled.
RAC	0	Report A Check Condition bit. <i>Ignored by the drive.</i>
SWP - Software Write Protect	0	This field controls the drive's write protection feature. If set, the drive returns check condition for the command with sense code/ASC/ASCQ - 07/27/00
RAERP	0	Ready AER Permission. <i>Not supported by the drive.</i>
UAAERP	0	Unit Attention AER Permission. <i>Not supported by the drive.</i>
EAERP	0	Error AER Permission. <i>Not supported by the drive.</i>
Ready AER Holdoff Period	0	<i>Not supported by the drive.</i>
Busy Timeout Period	FFFFh	<i>Not supported by the drive.</i>
Extended Self Test Completion Time	200h	Contains advisory data that indicates the approximate time, in seconds, required to complete an extended self-test when not interrupted by subsequent commands and no errors occur during execution of the self-test.

NOTES:

If the Queue Algorithm Modifier specifies restricted re-ordering (0000b), commands are not allowed to execute concurrently. If this field specifies unrestricted re-ordering (0001b), concurrent I/O execution is allowed.

5.12.11 Notch and Partition Page (0Ch)

The Notch and Partition page contains parameters for drives that implement a variable number of sectors per cylinder, and, support this page. Each section of the logical unit with a different number of sectors is referred to as a *notch* (or band). Each notch must span a set of consecutive logical blocks in the logical unit, the notches must not overlay, and no logical blocks can be excluded from the notches. The data format and field descriptions are described in Table 5-70 and Table 5-71 respectively.

Table 5-70 *Notch and Partition Page-(Page 0Ch)*

Byte	Bit	7	6	5	4	3	2	1	0
0		PS	Rsvd	Page Code (0Ch)					
1		Page Length (16h)							
2		ND	LPN	Reserved					
3		Reserved							
4 – 5		Maximum Number of Notches							
6 – 7		Active Notch							
8 – 11		Starting Boundary							
12 – 15		Ending Boundary							
16 – 23		Pages Notched							

Table 5-71 *Notch and Partition Page-Field Descriptions*

Data Field	Description
PS	Parameters Savable. This bit is only used with the MODE SENSE command. The returned value of 0 indicates that the target is not capable of saving this page in a non-volatile vendor specific location. This bit is reserved with the MODE SELECT command.
ND	Notched (banded) Drive. A value of 0 indicates that the drive is not notched and that all other parameters in this page are returned as 0. The default value of 1 indicates that the device is notched. For each supported active notch value, this page defines the starting and ending boundaries of the notch.
LPN	Logical or Physical Notch. The default value of 0 indicates that the boundaries of the notch are based on the physical parameters of the logical unit. The cylinder is considered most significant and the head least significant. A value of 1 indicates that the notch boundaries are based on logical blocks of the logical unit.
MaxNotch	This field indicates the maximum number of notches supported by the logical unit. The value cannot be changed.
Active Notch	The Active Notch field indicates the notch referenced by this, and subsequent MODE SELECT and MODE SENSE commands. The value in this field can be changed by a later MODE SELECT command. The value of the active notch must be = 0 and = 24. An active notch of 0 indicates that this and subsequent MODE SELECT and MODE SENSE commands refer to the parameters that apply across all notches.
Starting Boundary	This field indicates the beginning of the active notch or, if the active notch is 0, the beginning boundary of the logical unit. If the LPN value is 1, the starting boundary is a Logical Block Address. If the LPN value is 0, the three most significant bytes designate the starting cylinder number and the least significant byte is the starting head number. This field is unchangeable. When used with the MODE SELECT command, this field is ignored.
Ending Boundary	This field indicates the ending of the active notch, or if the active notch is zero, the ending of the logical unit. If the LPN value is 0, the 3 most significant bytes designate the ending cylinder and the least significant byte is the ending head number.
Pages Notched	This is a 64-bit bitmap that indicates which MODE command pages contain parameters that may be different for different notches. The most significant bit (MSB) corresponds to Page 3Fh (Vendor-Unique Caching Page), and the least significant bit (LSB) corresponds to Page 00h (Unit Attention Control Page). If a bit is 1, the corresponding mode page contains parameters that may be different for different notches. If a bit is 0, the corresponding mode page contains parameters that are constant for all notches. The value of 1 in bits 3 and 12 indicates that Page 03h and 0Ch contain different parameters for different notches.

5.12.12 Port Control Mode Page (19h)

The port control mode page (see Table 5-72 and Table 5-73) contains those parameters that affect SPI SCSI device port operation options. The page will be implemented by LUN 0 of all SPI SCSI devices. The page will not be implemented by logical units other than LUN 0. The implementation of any bit and its associated functions is optional. The page follows the MODE SENSE/MODE SELECT rules specified by SCSI Primary Commands-2 standard.

The target will maintain an independent set of port control mode page parameters for each initiator. The parameters saveable bit in the mode page format header returned with MODE SENSE command will be set to zero if the long mode page format is being used (LONG bit set to one), indicating the parameters are not saved through resets.

Table 5-72 *Port Control Page Short Format*

Bit Byte	7	6	5	4	3	2	1	0
0	PS	LONG(0)	Page Code (19h)					
1	Page Length (06h)							
2	Reserved				Protocol Identifier (1h)			
3	Reserved							
4	MSB	Synchronous Transfer Timeout						
5	Synchronous Transfer Timeout							LSB
6	Reserved							
7	Reserved							

The Protocol Identifier field of 1h indicates the protocol that this mode page applies to a SPI SCSI device. See the SCSI Primary Commands-2 standard for other port control page protocol identifiers.

The Synchronous Transfer Timeout field indicates the maximum amount of time in 1 millisecond increments that the target will wait before generating an error by doing an unexpected bus free. The target will only go to a BUS FREE phase if one of the following events causes the timer, once started, to not reset or reload before expiring:

- If there is a REQ transition when there are no outstanding REQs waiting for an ACK then load and start the timer.
- If there is a REQ transition when there are any outstanding REQs waiting for an ACK then there is no effect on the timer.
- If there is an ACK transition when there are outstanding REQs waiting for an ACK then load and start the timer.
- If after an ACK transition there are no outstanding REQs waiting for an ACK then stop the timer.
- A SYNCHRONOUS TRANSFER TIMEOUT field value of 0000h indicates that the function is disabled. A value of FFFFh indicates an unlimited period.

Table 5-73 *Port Control Page Long Format*

Bit Byte	7	6	5	4	3	2	1	0
0	PS	LONG (1)	Page Code (19h)					
1	Subpage Code							
2	MSB	Page Length (n-3)						
3	Page Length (n-3)							LSB
4	Reserved							
5	Reserved				Protocol Identifier (1h)			
6	Protocol Specific Mode Parameters							
n								

The Subpage Code field indicates which subpage is being accessed. Subpage Code values are listed in Table 5-93. If the Subpage Code field is zero, the target will return a CHECK CONDITION status. The sense key will be set to ILLEGAL REQUEST and the additional sense code set to ILLEGAL FIELD IN PARAMETER LIST.

The Page Length field specifies the length in bytes of the subpage protocol specific mode page parameters after the Page Length field.

The Protocol Identifier field of 1h indicates the protocol that this mode page applies to a SPI SCSI device. See the SCSI Primary Commands-2 standard for other port control page protocol identifiers.

5.12.12.1 Margin Control Subpage (Sub Page 1 of Mode Page 19)

The Margin Control subpage (see Table 5-74) contains parameters that set and report margin control values for usage between the initiator/target pair on subsequent synchronous and paced transfers.

A MODE SELECT command will return the current settings for the initiator/target pair. Fields that are not implemented will be reported as zero.

Table 5-74 Margin Control Subpage (01h)

Bit Byte	7	6	5	4	3	2	1	0
0	Reserved							
1	Driver Strength				Reserved			
2	Drive's Asymmetry				Driver Precompensation			
3	Drive's Slew Rate				Reserved			
4	Reserved							
5	Reserved							
6	Reserved							
7	Vendor Specific							
8	Reserved							
9	Reserved							
10	Reserved							
11	Reserved							
12	Reserved							
13	Reserved							
14	Reserved							
15	Reserved							

Table 5-75 Field Descriptions

Field	Description
Driver Strength	The DRIVER STRENGTH field indicates the relative amount of driver source currently used by the driver. The Driver Strength field affects both the strong and weak drivers. A larger value indicates more driver source current.
Driver Pre-compensation	The DRIVER PRECOMPENSATION field indicates the relative difference between the weak driver and the strong driver amplitudes when precompensation is enabled. A larger value indicates a larger difference between the weak and strong amplitudes.
Driver Asymmetry	The Driver Asymmetry field indicates the relative difference between the amplitudes of asserted and negated signals launched from the driver. A larger value indicates a relatively stronger asserted signal compared to the negated signal.
Driver Slew rate	The DRIVER SLEW RATE field indicates the relative difference between the assertion and negation magnitudes divided by the rise or fall time. A larger value indicates a faster slew rate.

The Driver Strength field indicates the relative amount of driver source currently used by the driver. The Driver Strength field affects both the strong and weak drivers. A larger value indicates more driver source current.

The Driver Precompensation field indicates the relative difference between the weak driver and the strong driver amplitudes when precompensation is enabled. A larger value indicates a larger difference between the weak and strong amplitudes.

The Driver Asymmetry field indicates the relative difference between the amplitudes of asserted and negated signals launched from the driver. A larger value indicates a relatively stronger asserted signal compared to the negated signal.

The Driver Slew Rate field indicates the relative difference between the assertion and negation magnitudes divided by the rise or fall time. A larger value indicates a faster slew rate.

5.12.12.2 Saved Training Configuration Values Subpage (Sub Page 2 of Mode Page 19)

The saved training configuration values subpage is used to report the SCSI device's saved training configuration values. These vendor specific values are maintained by the SCSI device when the retain training information option is enabled. The fields are listed in Table 5-76; however the content of the fields is vendor specific. Only values for the current I_T nexus are reported.

Table 5-76 *Saved Training Configuration Subpage (02h)*

Bit Byte	7	6	5	4	3	2	1	0
0	Reserved							
3	Reserved							
4	DB(0) Value							
7	DB(0) Value							
64	DB(15) Value							
67	DB(15) Value							
68	P_ORCA Value							
71	P_ORCA Value							
72	P1 Value							
75	P1 Value							
76	BSY Value							
79	BSY Value							
80	SEL Value							
83	SEL Value							
84	RST Value							
87	RST Value							
88	REQ Value							
91	REQ Value							
92	ACK Value							
95	ACK Value							
96	ATN Value							
99	ATN Value							
100	C/D Value							
103	C/D Value							
104	I/O Value							
107	I/O Value							
108	MSG Value							
111	MSG Value							
112	Reserved							
227	Reserved							

5.12.12.3 Negotiated Settings Subpage (Sub Page 3 of Mode Page 19)

The negotiated settings subpage, shown in Table 5-77, is used to report the negotiated settings of a target for the current I_T nexus.

Table 5-77 *Negotiated Settings Subpage (03h)*

Bit Byte	7	6	5	4	3	2	1	0
0	Transfer Period Factor							
1	Reserved							
2	Req/Ack Offset							
3	Transfer Width Exponent							
4	Reserved	Protocol Option Bits						
5	Reserved				Transceiver Mode	SENT_PCOMP_EN	RECEIVED_PCOMP_EN	
6	Reserved							
7	Vendor Specific							

The Transfer Period Factor field indicates the negotiated transfer period factor for the current I_T nexus.

Table 5-78 *Transceiver Period Factor*

Code	Transceiver Period Factor
0-7	Reserved
8	TP=6.25ns Fast - 160 (U320)
9	TP=12.5ns Fast - 80 (U160)
0x0A	TP=25ns Fast - 40 (Ultra2)
0x0B	TP=30.3 ns Fast - 40
0x0C-0x18	Fast - 20
0x19-0x31	Fast - 10
0x32-0xFF	Fast - 5

The Req/Ack Offset field indicates the negotiated REQ/ACK offset for the current I_T nexus.

The Transfer Width Exponent field indicates the negotiated transfer width exponent for the current I_T nexus.

The Protocol Options Bits field contain the negotiated protocol options for the current I_T nexus.

The TRANSCEIVER_MODE field specifies the current bus mode of the target as defined in Table 5-79.

Table 5-79 *Transceiver Mode*

Code	Transceiver Mode
00b	Unknown (device not capable of reporting bus mode)
01b	Single-ended
10b	Low-Voltage Differential
11b	High-Voltage Differential

The SENT PCOMP_EN bit contains the value of the PCOMP_EN bit sent by the target for the current I_T nexus.

The RECEIVED PCOMP_EN bit contains the value of the PCOMP_EN bit received by the target for the current I_T nexus.

5.12.12.4 Report Transfer Capabilities Subpage - (Sub Page 4 of Mode Page 19)

The report transfer capabilities subpage is used to report the transfer capabilities for the SCSI target port. The values in this subpage are not changeable via a MODE SELECT command.

Table 5-80 Report Transfer Capabilities Subpage (4)

Bit Byte	7	6	5	4	3	2	1	0
0	Minimum Transfer Period Factor							
1	Reserved							
2	Maximum Req/Ack Offset							
3	Maximum Transfer Width Exponent							
4	Protocol Option Bits Supported							
5-7	Reserved							

Table 5-81 Field Descriptions

Field	Description																											
Minimum Transfer Period Factor	The MINIMUM TRANSFER PERIOD FACTOR field shall be set to the smallest value of the transfer period factor																											
Maximum Req/ack Offset	The MAXIMUM REQ/ACK OFFSET shall be set to the largest value of the REQ/ACK offset																											
Maximum Transfer Width Exponent	The MAXIMUM TRANSFER WIDTH EXPONENT shall be set to the largest value of the transfer width exponent																											
Protocol Options Bits Supported	<p>The SCSI target port shall set the bits in the PROTOCOL OPTIONS BITS SUPPORTED field to indicate the protocol options supported by the SCSI target port.</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Name</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7</td> <td>PCOMP_EN</td> <td>Precompensation enable</td> </tr> <tr> <td>6</td> <td>RTI</td> <td>Retain training information</td> </tr> <tr> <td>5</td> <td>RD_STRM</td> <td>Read streaming and read flow control enable</td> </tr> <tr> <td>4</td> <td>WR_FLOW</td> <td>Write flow control enable</td> </tr> <tr> <td>3</td> <td>HOLD_MCS</td> <td>Hold margin control settings</td> </tr> <tr> <td>2</td> <td>QAS_REQ</td> <td>QAS enable request</td> </tr> <tr> <td>1</td> <td>DT_REQ</td> <td>DT clocking enable request</td> </tr> <tr> <td>0</td> <td>IU_REQ</td> <td>Information units enable request</td> </tr> </tbody> </table>	Bit	Name	Description	7	PCOMP_EN	Precompensation enable	6	RTI	Retain training information	5	RD_STRM	Read streaming and read flow control enable	4	WR_FLOW	Write flow control enable	3	HOLD_MCS	Hold margin control settings	2	QAS_REQ	QAS enable request	1	DT_REQ	DT clocking enable request	0	IU_REQ	Information units enable request
Bit	Name	Description																										
7	PCOMP_EN	Precompensation enable																										
6	RTI	Retain training information																										
5	RD_STRM	Read streaming and read flow control enable																										
4	WR_FLOW	Write flow control enable																										
3	HOLD_MCS	Hold margin control settings																										
2	QAS_REQ	QAS enable request																										
1	DT_REQ	DT clocking enable request																										
0	IU_REQ	Information units enable request																										

5.12.13 Information Exceptions Control Page (1Ch)

This page defines the methods used by the disk drive to control the reporting and operations of specific informational exception conditions. This page applies only to informational exceptions that report an additional sense code of FAILURE PREDICTION THRESHOLD EXCEEDED to the application client.

Informational exception conditions occur as the result of vendor-specific events within a disk drive. An informational exception condition may occur asynchronous to any commands issued by an application client.

The data format and field descriptions are described in Table 5-82 and Table 5-83 respectively. The Codes Used by the MRIE Field are described in Table 5-84.

Table 5-82 *Information Exceptions Control Page-(Page 1Ch)*

Bit Byte	7	6	5	4	3	2	1	0
0	PS	Page Code (1Ch)						
1	Page Length (0Ah)							
2	Perf	Reserved	EWASC	DExcpt	Test	Rsv'd	LogErr	
3	Reserved				MRIE			
4 - 7	(MSB)	Interval Timer						(LSB)
8 - 11	(MSB)	Report Count						(LSB)

Table 5-83 *Information Exceptions Control Page-Field Descriptions (Page 1Ch)*

Field	Default Value	Description
PS	1	Parameters Savable. This bit is only used with the MODE SENSE command. A returned value of 1 indicates that the target is capable of saving the page in a non-volatile, vendor-specific location. This bit is reserved with the MODE SELECT command.
PERF	1	Performance. This bit enables/disables performance-sensitive functions. For Atlas 10K V, this means enabling or disabling flushing of SMART data to the media (done hourly). The prediction algorithm is designed to reset the current accumulated data approximately every 100,000 blocks read or written per head. When this bit = 0, current accumulated SMART data is preserved across power cycles. When this bit = 1, current accumulated SMART data is lost across power cycles.
EWASC	1	Enable Warning ASC. When this bit = 0, reporting of thermal warning is disabled. When this bit = 1, thermal warning reports are enabled.
DExcpt	0	Disable Exceptions. This bit is used to enable or disable SMART exception reporting. When this bit = 0, SMART exception reporting is enabled. When this bit = 1, all SMART functions are disabled, including reporting SMART exceptions, running the prediction algorithm, and accumulating SMART data.
Test	0	When this bit = 0, the test function is disabled. When this bit = 1, the disk drive reports false SMART trip reports as prescribed in the MRIE field (see below). Interval timer field and report count and method fields apply. When the test bit forces a SMART failure, the failure is reported with a sense code/qualifier of 5D/FF.
LogErr	0	When this bit = 0, it indicates that log sense information is vendor-unique. NOTE: LogErr = 0 is the only value supported.
MRIE	4	Method of Reporting Informational Exceptions. The value within this field indicates the method that will be used by the disk drive to report informational exception conditions, as shown in Table 5-84. The relative priority of reporting of multiple informational exceptions is vendor-specific.
Interval Timer	0	This field indicates the period, in 100 millisecond increments, for reporting that an informational exception condition has occurred. The disk drive will not report such conditions more frequently than the time specified by the Interval Timer field and as soon as possible after the timer interval has elapsed. After the condition has been reported, the interval timer restarts. A value of 0 in the Interval Timer field indicates that the disk drive will only report the informational exception condition once. A value of FFFFFFFFh in the Interval Timer field indicates that the timer interval is vendor-specific. If the disk drive's predictor detects a SMART threshold trip, it reports the failure only if it has not reported a failure within the last interval timer period. A value of 0 means that there is no required waiting between SMART reports. NOTE: The Interval Timer field does not apply to MRIE mode 6.
Report Count	1	The value contained in this field indicates the number of times to report an informational exception condition to the application client. A value of 0 in this field indicates there is no limit on the number of times the disk drive will report an informational exception condition. Each time the disk drive detects a SMART failure, the SMART status is reported the number of times indicated by the value of the Report Count field. The disk drive then resets its cumulative values and continues with its predicting activity. The SMART threshold trip remains permanent (or until a LOG SELECT command clears it) and can be retrieved at any time using MRIE mode 6. NOTE: The Report Count field does not apply to MRIE mode 6.

Table 5-84 *Codes Used by the MRIE Field*

Field	Description
0h	No reporting of informational exception condition. This method instructs the disk drive not to report these conditions.
1h	Asynchronous Event Reporting. This method of reporting instructs the disk drive to report informational exception conditions by using the rules for asynchronous event reporting as described in the SCSI-3 Architecture Model and the relevant Protocol Standard. The sense key will be set to RECOVERED ERROR and the additional sense code will indicate the cause of the informational exception condition.
2h	Generate Unit Attention. This method of reporting instructs the disk drive to report informational exception conditions by returning a CHECK CONDITION status on any command. The sense key is set to UNIT ATTENTION and the additional sense code will indicate the cause of the information exception condition. The command that has the CHECK CONDITION will not be executed before the informational exception condition is reported.
3h	Conditionally Generate Recovered Error. This method of reporting instructs the disk drive to report informational exception conditions dependent on the value of the PER bit of the Error Recovery Parameters Mode Page, by returning a CHECK CONDITION status on any command. The sense key will be set to RECOVERED ERROR and the additional sense code will indicate the cause of the informational exception condition. The command that has the CHECK CONDITION will complete before and informational exception condition is reported.
4h	Unconditionally Generate Recovered Error. This method of reporting instructs the disk drive to report informational exception conditions regardless of the value of the PER bit of the Error Recovery Parameters Mode Page, by returning a CHECK CONDITION status on any command. The sense key will be set to RECOVERED ERROR and the additional sense code will indicate the cause of the informational exception condition. The command that has the CHECK CONDITION will complete without error before any informational exception condition will be reported.
5h	Generate No Sense. This method of reporting instructs the disk drive to report informational exception conditions by returning a CHECK CONDITION status on any command. The sense key will be set to NO SENSE and the additional sense code will indicate the cause of the informational exception condition. The command that has the CHECK CONDITION will complete without error before any informational exception condition is reported.
6h	Only Report Informational Exception Condition On Request. This method of reporting instructs the disk drive to preserve the informational exception(s) data. To learn about these exception conditions, the application client polls the disk drive via an unsolicited REQUEST SENSE command. The sense key is set to NO SENSE and the additional sense code will indicate the cause of the informational exception condition.
7h – Bh	Reserved.
Ch – Fh	Vendor-specific.

5.13 MODE SELECT (10) Command (55h)

The MODE SELECT (10) command allows the initiator to modify certain mode pages and thereby control some of the drive's operational characteristics. Refer to MODE SELECT (6) for additional descriptions. See Table 5-85 for the data format of the MODE SELECT (10) Command Descriptor Block. See MODE SELECT (6), Table 5-88, for descriptions of the fields. See Table 5-86 for the Mode Parameter Header (10 Byte) — Data Format and Table 5-87 for the Mode Parameter Block Descriptor — Data Format. See Table 5-48 for mode parameter block field descriptions.

Table 5-85 *MODE SELECT (10) Command Descriptor Block-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (55h)							
1	Reserved			PF	Reserved			SP
2 - 6	Reserved							
7 - 8	Parameter List Length							
9	Control							

Table 5-86 *Mode Parameter Header (10-Byte)-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0 - 1	Mode Data Length							(MSB)
	(LSB)							
2	Medium Type							
3	Device-Specific Parameter							
4 - 5	Reserved							
6 - 7	Block Descriptor Length							(MSB)
	(LSB)							

Table 5-87 *Mode Parameter Block Descriptor-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0 - 3	Number of Blocks							(MSB)
	(LSB)							
4	Density Code							
5 - 7	Block Length							(MSB)
	(LSB)							

Table 5-88 *Mode Parameter Block Descriptor-Field Descriptions*

Data Field	Description
Number of Blocks	The value in this field specifies the number of logical blocks on the disk to which the density code and block length fields apply. A value of 0 in this field indicates that all of the remaining logical blocks on the disk have the medium characteristics specified.
Density Code	The Density Code field is unique for each device type. Direct access devices value is zero.
Block Length	The value in this field specifies the length in bytes of each logical block described by the block descriptor.

5.14 MODE SENSE (6) Command (1Ah)

The MODE SENSE (6) command allows the initiator to read the drive's sector descriptor and mode pages. The MODE SENSE Command Descriptor Sector is shown in Table 5-89 and the data fields are described in Table 5-90. The Mode Parameter Header is shown in Table 5-91 and the data fields are described in Table 5-92.

Table 5-89 *MODE SENSE (6) Command Descriptor Block-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (1Ah)							
1	Reserved				DBD		Reserved	
2	PC		Page Code					
3	Sub Page Code							
4	Allocation Length							
5	Control							

Table 5-90 *MODE SENSE Command-Field Descriptions*

Data Field	Description
DBD	Disable Block Descriptor. When the value is 0, the drive returns its Block Descriptor followed by the requested page(s). When the value is 1, the drive returns only the requested pages.
PC	Page Control. Determines which type of page values are returned. These values are: 00b Current 01b Changeable 10b Default 11b Saved
Page Code	The Page Code field determines the pages to be returned to the initiator in the command's data-in buffer transfer. Refer to Table 5-49 and Table 5-93 for a complete list of the Mode Pages (changeable and read-only). A Page Code of 3Fh causes all supported pages to be returned.
Sub Page Code	The Sub Page Code field determines the sub pages to be returned to the initiator in the command's data-in buffer transfer. A Sub Page Code of 0xFF causes all supported sub pages to be returned.

Table 5-91 *Mode Parameter Header (6 Byte)-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0	Mode Data Length							
1	Medium Type							
2	Device-Specific Parameter							
3	Block Descriptor Length							

Table 5-92 Mode Parameter Header and Block Descriptor-Field Descriptions

Data Field	Description
Mode Data Length	Length of the returned data
Medium Type	Always zero (0)
Device-Specific Parameter	Bit 7 - This bit is set if the drive is write protected Bit 4 - This bit is always set to indicate that the drive supports cache memory.
Block Descriptor	Zero (0) if Block Descriptor (DBD) in MODE SENSE Command Descriptor Length sector is 0. If DBD is 1, then the Block Descriptor Length is eight (8).
Number of Blocks	The value in this field specifies the number of logical blocks on the unit that are affected by the current MODE SELECT command. When this field = 0, the MODE SELECT parameters are put into effect for all logical blocks on the unit. The maximum value allowed is the total number of initiator-addressable logical blocks for the value specified in the Block Length field. A value that is less than the maximum value will decrease the number of initiator-accessible logical blocks to the amount specified.
Block Length	The value in this field specifies the length, in bytes, of each block.

5.14.1 Read-Only Mode Pages

Table 5-93 lists read-only mode pages. Table 5-49 lists initiator-changeable mode pages.

Table 5-93 Read-Only Mode Pages

Page Code (Sub Page)	Page Name	Reports the drive's	Size (bytes)
03h	Format Device	Medium format	24
04h	Rigid Disk Geometry	Total cylinder head and rotational speed	24

5.14.2 Format Device Page (03h)

The Format Drive Page (Table 5-94) describes the drive's medium format. The data fields are described in Table 5-95. Note that the page is *savable* but no parameters are *changeable*.

Table 5-94 Format Device Page-(Page 3)

Byte	Bit	7	6	5	4	3	2	1	0
0		PS	Rsvd	Page Code (03h)					
1		Page Length (16h)							
2 - 3		Tracks Per Zone							
4 - 5		Alternate Sectors/Zone							
6 - 7		Alternate Tracks/Zone							
8 - 9		Alternate Tracks/Logical Unit							
10 - 11		Sectors/Track							
12 - 13		Data Bytes/Physical Sector							
14 - 15		Interleave							
16 - 17		Track Skew Factor							
18 - 19		Cylinder Skew Factor							
20		SSEC	HSEC	RMB	SURF	Reserved			

21 - 23	Reserved
---------	----------

Table 5-95 *Format Device Page-Field Descriptions (Page 3)*

Field	Description
PS	Parameters Savable. This bit is only used with the MODE SENSE command. This bit is reserved with the MODE SELECT command. A PS bit value of 1 indicates that the disk drive is capable of saving the page in a non-volatile vendor specific location.
Tracks/Zone	The number of physical tracks per cylinder.
Alternate Sectors/Zone	The number of spare locations within a notch.
Alternate Tracks/Zone	The drive does not have any alternate tracks.
Alternate Tracks/Logical Unit	The number of tracks reserved for defect handling.
Sectors/Track	The number of Physical Block Address sectors per track. The value depends on the Notch value selected in the Notch and Partitions page.
Data Bytes/Physical Sector	This value is the current disk sector size.
Interleave	The drive does not support interleave.
Track Skew Factor	The number of physical sectors between the last logical sector on one track and the first logical sector on the next track within a cylinder.
Cylinder Skew Factor	The number of physical sectors between the last logical sector on one cylinder and the first logical sector on the next cylinder.
SSEC	The drive does not use soft sector formatting.
HSEC	The drive supports hard sector formatting only.
RMB	The drive does not support removable logical media.
SURF	A Surface bit value of 0 indicates that the drive allocates progressive addresses to all logical blocks within a cylinder prior to allocating addresses on the next cylinder.

5.14.3 Rigid Disk Geometry Page (04h)

The Rigid Disk Geometry Page describes the drive geometry. The data format and field descriptions are described in Table 5-96 and Table 5-97 respectively.

Table 5-96 *Rigid Disk Geometry Page-(Page 4)*

Byte	Bit	7	6	5	4	3	2	1	0
0		PS	Rsvd	Page Code (04h)					
1		Page Length (16h)							
2 - 4		Number of Cylinders							
5		Number of Heads							
6 - 8		Starting Cylinder - Write Precompensation							
9 - 11		Starting Cylinder - Reduced Write Current							
12 - 13		Drive Step Rate (0)							
14 - 16		Landing Zone Cylinder							
17		Reserved						RPL	
18		Rotational Offset							
19		Reserved							
20 - 21		Medium Rotation Rate							
22 - 23		Reserved							

Table 5-97 Rigid Disk Geometry Page-Field Descriptions (Page 4)

Field	Description
PS	Parameters Savable. This bit is only used with the MODE SENSE command. A returned value of 1 indicates that the target is capable of saving the page in a non-volatile, vendor-specific location. The bit is reserved with the MODE SELECT command.
Number of Cylinders	The number of physical cylinders that hold host data.
Number of Heads	The number of heads that can access host data.
Starting Cylinder	The value in this field is the physical cylinder at
Write Precompensation	which write precompensation is to begin. The first cylinder number is 0. If the starting cylinder is equal to the value in the Number of Cylinders field, write precompensation is disabled.
Starting Cylinder Reduced Write Current	The value in this field is the physical cylinder at which write current is reduced. The first cylinder number is 0. If the starting cylinder for reduced write current is equal to the value in the Number of Cylinders field, reduced write current is disabled.
Drive Step Rate	Applies to stepper-motor actuators. Not applicable.
Landing Zone Cylinder	This field indicates a two's complement location where the drive positions the disk heads. A negative value indicates that the heads are positioned below cylinder 0 by that number of cylinders. A positive value greater than the number of cylinders indicates that the heads are positioned beyond the cylinders used for data storage at the location specified.
RPL	Rotational Position Locking is used for spindle synchronization in storage arrays. The modes are: 00b – Disabled 10b – Master 01b – Slave 11b – Not supported RPL can be overridden if the Maxtor (Vendor) Special Function Control page Auto field value = 1. <i>This drive does not support spindle synchronization.</i>
Rotational Offset	Rotational Offset indicates the amount of rotational offset that the drive uses when synchronized. The rotational offset is applied in the retarded direction (lagging the synchronized spindle master). The value in the field is determined by the following formula: $\text{Offset} = \frac{0 \dots 255}{256}$ For example, a numerator of 128 indicates an offset of 1/2 revolution (180°). The valid parameter values are 0 to 255. The specified value may be rounded. Rotational offset is not used when the drive is configured as the synchronized-spindle master.
Medium Rotation Rate	The disk rotation speed in RPM.

5.15 MODE SENSE (10) COMMAND (5Ah)

The MODE SENSE command allows the initiator to read the drive's sector descriptor and mode pages.

See Table 5-98 for MODE SENSE (10) Command Descriptor Block – Data Format, Table 5-99 for Mode Parameter Header (10 Byte) — Data Format, and Table 5-100 for Mode Parameter Block Descriptor – Data Format. Refer to Table 5-90, Table 5-92, and Table 5-101 for descriptions of the fields.

Table 5-98 *MODE SENSE (10) Command Descriptor Block-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (5Ah)							
1	Reserved				DBD	Reserved		
2	PC		Page Code					
3	Sub Page Code							
4 – 6	Reserved							
7 – 8	(MSB) Allocation Length							(LSB)
9	Control							

Table 5-99 *Mode Parameter Header (10 Byte)-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0 - 1	(MSB) Mode Data Length							(LSB)
2	Medium Type							
3	Device-Specific Parameter							
4 - 5	Reserved							
6 - 7	Block Descriptor Length							

Field Descriptions for the Mode Parameter Header are given in the table that follows.

Table 5-100 *Mode Parameter Block Descriptor-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0 - 3	(MSB) Number of Blocks							(LSB)
4	Density Code							
5 - 7	(MSB) Block Length							(LSB)

Table 5-101 *Mode Parameter Header & Block Descriptor-Field Descriptions*

Data Field	Description
Mode Data Length	Length of the returned data.
Medium Type	Always zero (0).
Device Specific Parameter	Bit 7 - This bit is set if the drive is write protected. Bit 4 - This bit is always set to indicate that the drive supports cache memory.
Block Descriptor Length	The value = 0 if no Block Descriptor is supplied; the length is 8 if a block descriptor is supplied.
Density Code	Ignored by the drive.
Number of Blocks	The value in this field specifies the number of logical blocks on the unit that are affected by the current MODE SELECT command. When this field = 0, the MODE SELECT parameters are put into effect for all logical blocks on the unit. The maximum value allowed is the total number of initiator-addressable logical blocks for the value specified in the Block Length field. A value that is less than the maximum value will decrease the number of initiator-accessible logical blocks to the amount specified.
Block Length	The value in this field specifies the length, in bytes, of each logical block.

5.16 PERSISTENT RESERVATION IN Command (5Eh)

The PERSISTENT RESERVATION IN command is a 10-byte command used to obtain information about persistent reservations and reservations that are active within a device server. It is used in conjunction with the PERSISTENT RESERVATION OUT command.

Note that a if device that has RESERVE (10), RESERVE (6), extent reservations, or SCSI Media Changer (SMC) element reservations active receives a PERSISTENT RESERVATION IN command, the PERSISTENT RESERVATION IN command is rejected with a RESERVATION CONFLICT status reported.

Table 5-102 illustrates the format of the PERSISTENT RESERVATION IN command. Table 5-103 explains the data fields of the command.

Table 5-102 PERSISTENT RESERVATION IN Command Descriptor Block-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (5Eh)							
1	Reserved			Service Action				
2 - 6	Reserved							
7 - 8	(MSB) Allocation Length (LSB)							
9	Control							

Table 5-103 PERSISTENT RESERVATION IN Command-Field Descriptions

Data Field	Description
Service Action	<p>Service actions that require information about persistent reservation and registrations may require enabling of nonvolatile memory within the logical unit. Service action codes available are:</p> <p>Code Name Description 00h Read Keys Reads all registered Reservation keys 01h Read Reservations Reads all current persistent reservations 02 – 1Fh Reserved</p> <p>A “Read Keys” service action requests that the device server return a parameter list that includes a header and a complete list of all of the reservation keys currently registered with the device server. If multiple initiators have registered with the same key, then the key is listed multiple times, once for each registration. Refer to Table 5-104 and Table 5-105 for information about Read Keys parameter data.</p> <p>A “Read Reservation” service action requests that the device server return a parameter list that contains a header and a complete list of all persistent reservations that are presently active in the device server and its extents. Note that duplicate persistent reservations from the same initiator are not reported. Refer to Table 5-108 and Table 5-107 for information about Read Reservations parameter data.</p>
Allocation Length	<p>This field indicates how much space has been reserved for the returned parameter list (Read Keys or Read Reservations parameters). The actual length of the parameter data is indicated in the parameter data field for those parameters.</p> <p>If the Allocation Length is not sufficient to contain the entire list of parameters, the first portion of the list that does fit is returned. If it is determined that the remainder of the list is required, the client should send a new PERSISTENT RESERVATION IN command with an Allocation Length field large enough to contain the entire list of parameters.</p>

The data format and field descriptions for the Read Keys Parameters are described in Table 5-104 and Table 5-105 respectively.

Table 5-104 Read Keys Parameters-Data Format

Byte	Bit	7	6	5	4	3	2	1	0
0 - 3	(MSB)	Generation							(LSB)
4 - 7	(MSB)	Additional Length ($n - 7$)							(LSB)
(Reservation Key List Follows in Bytes 8 – n)									
8 - 15	(MSB)	First Reservation Key							(LSB)
$n - 7$ to n	(MSB)	Last Reservation Key							(LSB)

Table 5-105 *Read Keys Parameters-Field Descriptions*

Data Field	Description
Generation	<p>The value in this field is a 32-bit counter in the device server that is incremented each time a PERSISTENT RESERVATION OUT command requests a Register, Clear, Pre-empt, or Pre-empt and Clear operation. Note that PERSISTENT RESERVATION IN commands do not increment the counter, nor do PERSISTENT RESERVATION OUT commands that perform a Reserve or Release service action, or by a PERSISTENT RESERVATION OUT command that is not done due to an error or a reservation conflict. The value in the Generation field is set to 0 as part of the power on or reset processes.</p> <p>The value in the Generation field allows the application client that examines the value to verify that the configuration of the initiators attached to a logical unit has not been modified by another application client without any notification of the application client doing the examination.</p>
Additional Length	<p>This field contains the count of the number of bytes that are in the Reservation Key list (bytes 8 – n). Note that this field contains the number of bytes in the reservation key list regardless of the value prescribed by the Allocation Length field in the command's CDB.</p>
Reservation Keys	<p>Each of the Reservation Keys appear as items in a list as bytes 8 through n. Each entry reflects an 8-byte reservation key registered with the device server via the PERSISTENT RESERVATION OUT, Reserve, Pre-empt, Pre-empt and Clear, or Register service actions. Each key can be examined by the application client for correlation with a set of initiators and SCSI ports</p>

The data format and field descriptions for the Read Reservations Parameters are described in Table 5-106 and Table 5-107 respectively.

Table 5-106 *Read Reservations Parameters-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0 - 3	(MSB) Generation (LSB)							
4 - 7	(MSB) Additional Length ($n - 7$) (LSB)							
8 - n	(MSB) Reservation Descriptors (See Figure 5-53 / Table 5-55 for detail of a Descriptor Field)(LSB)							

Table 5-107 *Read Reservations Parameters-Field Descriptions*

Data Field	Description
Generation	<p>The value in this field is a 32-bit counter in the device server that is incremented each time a PERSISTENT RESERVATION OUT command requests a Register, Clear, Pre-empt, or Pre-empt and Clear operation. Note that PERSISTENT RESERVATION IN commands do not increment the counter, nor do PERSISTENT RESERVATION OUT commands that perform a Reserve or Release service action, or by a PERSISTENT RESERVATION OUT command that is not done due to an error or a reservation conflict. The value in the Generation field is set to 0 as part of the power on or reset processes.</p> <p>The value in the Generation field allows the application client that examines the value to verify that the configuration of the initiators attached to a logical unit has not been modified by another application client without any notification of the application client doing the examination.</p>
Additional Length	This field contains the count of the number of bytes of Reservation descriptors (bytes $8 - n$). Note that this field contains the number of bytes regardless of the value prescribed by the Allocation Length field in the command's CDB.
Reservations Descriptors	One Reservation descriptor is reported for each unique persistent reservation on the logical unit when the PERSISTENT RESERVATION IN command has indicated a Read Reservations action. Table 5-112 and Table 5-109 detail the contents of each Reservation Descriptors field.

The data format and field descriptions for the PERSISTENT RESERVATION IN Read Reservations Descriptor are described in Table 5-108 and Table 5-109 respectively.

Table 5-108 *PERSISTENT RESERVATION IN Read Reservations Descriptor-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0 - 7	(MSB) Reservation Key							(LSB)
8 - 11	(MSB) Scope-Specific Address							(LSB)
12	Reserved							
13	Scope				Type			
14 - 15	(MSB) Extent Length							(LSB)

Table 5-109 *PERSISTENT RESERVATION IN Read Reservations Descriptor-Field Descriptions*

Data Field	Description
Reserva- tion Key	The Reservation key field contains an 8-byte value that identifies the reservation key under which the per- sistent reservation is held.
Scope- Specific Address	If the Scope field (Byte 13, bits 4 – 7) represents an “Extent” reservation, the Scope-Specific Address field contains the logical block address (LBA) of the first block of the extent and the Extent Length field (Bytes 14 – 15) contains the number of blocks in the extent. If the Scope field represents an “Element” reservation, the Scope-Specific Address field contains the Ele- ment address, zero-filled in the most significant bytes to fit the field and the Extent Length field is set to zero.
Scope	The value in this field indicates whether a persistent reservation applies to an entire logical unit, to a part of the logical unit (defined as an extent), or to an element. The values for the Scope field are: Code Name Description 0h LU Logical Unit. Persistent reservation applies to the full logical unit. The LU scope is therefore imple- mented by all device servers that implement PERSISTENT RESERVATION OUT. 1h Extent Persistent reservation applies to the specified extent. When Extent is the scope, it indicates that the persistent reservation applies to the extent of the logical unit defined by the Scope-Specific Address and Extent Length fields in the PERSISTENT RESERVATION OUT command’s parameter list. Note that an extent is defined only for devices defining contiguous logical block addresses. The Extent scope is optional for all device servers that implement PERSISTENT RESERVATION OUT. 2h Element Persistent reservation applies to the specified element. When Element is the scope, it indicates that the persistent reservation applies to the element of the logical unit defined by the Scope-Specific Ad- dress field in the PERSISTENT RESERVATION OUT parameter list. An element is as defined by the SCSI-3 Medium Changer Commands (SMC) standard. The Element scope is optional for all device servers that implement PERSISTENT RESERVATION OUT. 3h – Fh Reserved
Type	The value of the Type field specifies the characteristics of the persistent reservation being established for all data blocks within the extent or within the logical unit. Refer to Table 5-56 for the applicable Type codes and their meanings

Table 5-110 presents the definitions of the characters of the available “Type” values from the Type field of the PERSISTENT RESERVATION IN Read Reservations parameters.

Each of the codes provides handling instructions for READ operations, for WRITE operations, and for subsequent attempts to establish persistent reservations (referred to as “Additional Reservations Allowed” in the table).

Table 5-110 *Persistent Reservation Type Codes and Their Meanings*

Code	Name	Description
0h	READ Shared	<p>READS: Shared; any application client on any initiator may execute commands that perform transfers from the disk to the initiator.</p> <p>WRITES: Prohibited; any command from any initiator that performs a transfer from the initiator to the disk results in a reservation conflict.</p> <p>ADDITIONAL RESERVATIONS: Allowed; any initiator may reserve the logical unit, extents, or elements as long as the persistent reservations do not conflict with any reservations already known to the device server.</p>
1h	WRITE Exclusive	<p>READS: Shared; any application client on any initiator may execute commands that perform transfers from the disk to the initiator.</p> <p>WRITES: Exclusive; any command from any initiator other than the initiator that holds the persistent reservation that attempts a transfer to the disk results in a reservation conflict.</p> <p>ADDITIONAL RESERVATIONS: Allowed; any initiator may reserve the logical unit, extents, or elements as long as the persistent reservations do not conflict with any reservations already known to the device server.</p>
2h	READ Exclusive	<p>READS: Exclusive; any command from any initiator other than the initiator that holds the persistent reservation that attempts a transfer from the disk results in a reservation conflict.</p> <p>WRITES: Shared; any application client on any initiator may execute commands that perform transfers to the disk from the initiator.</p> <p>ADDITIONAL RESERVATIONS: Allowed; any initiator may reserve the logical unit, extents, or elements as long as the persistent reservations do not conflict with any reservations already known to the device server.</p>
3h	Exclusive Access	<p>READS: Exclusive; any command from any initiator other than the initiator holding the persistent reservation that attempts a transfer from the disk results in a reservation conflict.</p> <p>WRITES: Exclusive; any command from any initiator other than the initiator holding the persistent reservation that attempts a transfer to the disk results in a reservation conflict.</p> <p>ADDITIONAL RESERVATIONS: Restricted; any PERSISTENT RESERVATION OUT command with the Reserve service action from any initiator other than the initiator holding the persistent reservation results in a reservation conflict. The initiator that holds the persistent reservation can reserve the logical unit, extents, or elements as long as the persistent reservations do not conflict with any reservations already known to the device server.</p>
4h	Shared Access	<p>READS: Shared; any application client on any initiator may execute commands that perform transfers from the disk to the initiator.</p> <p>WRITES: Shared; any application client on any initiator may execute commands that perform transfers to the disk from the initiator.</p> <p>ADDITIONAL RESERVATIONS: Restricted; any PERSISTENT RESERVATION OUT command with the Reserve service action from any initiator other than the initiator holding the persistent reservation results in a reservation conflict. The initiator that holds the persistent reservation can reserve the logical unit, extents, or elements as long as the persistent reservations do not conflict with any reservations already known to the device server.</p>

Code	Name	Description
5h	WRITE Exclusive, Registrants Only	<p>READS: Shared; any application client on any initiator may execute commands that perform transfers from the disk to the initiator.</p> <p>WRITES: Exclusive; any command from an initiator that has not previously performed a Register service action with the device server that attempts a transfer to the disk results in a reservation conflict</p> <p>ADDITIONAL RESERVATIONS: Allowed; any initiator may reserve the logical unit, extents, or elements as long as the persistent reservations to not conflict with any reservations already known to the device server.</p>
6h	Exclusive Access, Registrants Only	<p>READS: Exclusive; any command from an initiator that has not previously performed a Register service action with the device server that attempts a transfer from the disk results in a reservation conflict.</p> <p>WRITES: Exclusive; any command from an initiator that has not previously performed a Register service action with the device server that attempts a transfer to the disk results in a reservation conflict</p> <p>ADDITIONAL RESERVATIONS: Allowed; any initiator may reserve the logical unit, extents, or elements as long as the persistent reservations to not conflict with any reservations already known to the device server.</p>
7h - Fh	Reserved	N/A

Table 5-111 represents graphically all possible combinations of Persistent Reservations Being Attempted when Persistent Reservations are already held for each of the types of persistent reservations.

Table 5-111 *When Do Conflicts Between Existing Reservations and New Reservations Exist?*

Persistent Reservation Being Attempted	Persistent Reservation Already Held						
	READ Shared LU EX	WRITE Exclusive LU EX	READ Exclusive LU EX	Exclusive Access * LU EX	Shared Access * LU EX	WRITE Exclusive RO LU EX	Exclusive Access RO LU EX
LU READ Shared EX	NN	YY	YY	YY	NN	YY	YY
	NN	YO	YO	YO	NN	YO	YO
LU WRITE Exclusive EX	YY	YY	YY	YY	YY	YY	YY
	YO	YO	YO	YO	YO	YO	YO
LU READ Exclusive EX	YY	YY	YY	YY	YY	YY	YY
	YO	YO	YO	YO	YO	YO	YO
LU Exclusive Access * EX	YY	YY	YY	YY	YY	YY	YY
	YO	YO	YO	YO	YO	YO	YO
LU Shared Access * EX	NN	YY	YY	YY	NN	YY	YY
	NN	YO	YO	YO	NN	YO	YO
LU WRITE Exclusive RO EX	YY	YY	YY	YY	YY	NN	NN
	YO	YO	YO	YO	OO	NN	NN
LU Exclusive Access RO EX	YY	YY	YY	YY	YY	NN	NN
	YO	YO	YO	YO	YO	NN	NN
KEY: LU = Logical Unit Scope EX = Extent or Element Scope RO = Registrants Only * = Conflicts with all reservation requests from other initiators.				N = No Conflict Y = Conflict O = Conflict occurs if extent or element overlaps with existing extent or element reservation.			

5.17 PERSISTENT RESERVATION OUT Command (5Fh)

The PERSISTENT RESERVATION OUT command is a 10-byte command used to reserve a logical unit or an extent within a logical unit for the exclusive or shared use by an initiator. The command is used in conjunction with the PERSISTENT RESERVATION IN command; it is not used with the RESERVE and RELEASE commands.

Persistent reservations conflict with reservations made via the RESERVE command. Initiators that perform PERSISTENT RESERVATION OUT actions are identified by a reservation key assigned by the application client. The client may use the PERSISTENT RESERVATION IN command to identify which other initiators within a system hold conflicting or invalid persistent reservations and use the PERSISTENT RESERVATION OUT command to preempt those reservations if necessary.

Note that since persistent reservations are not reset by the TARGET RESET task management function or other global actions, they can be used to enact device sharing among multiple initiators. The PERSISTENT RESERVATION OUT and PERSISTENT RESERVATION IN commands provide the means for resolving contentions in multiple-initiator systems with multiple port target. By using the reservation key to identify persistent reservations, it is possible to determine which ports hold conflicting persistent reservations and to take over such reservations from failing or “greedy” initiators.

illustrates the format of the PERSISTENT RESERVATION OUT command. Table 5-113 explains the data fields of the command.

Table 5-112 PERSISTENT RESERVATION OUT Command Descriptor Block-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (5Fh)							
1	Reserved			Service Action				
2	Scope				Type			
3 - 6	Reserved							
7 - 8	(MSB) Parameter List Length (18h) (LSB)							
9	Control							

Table 5-113 PERSISTENT RESERVATION OUT Command-Field Descriptions

Data Field	Description		
Service Action	Service actions that require information about persistent reservation and registrations may require enabling of nonvolatile memory within the logical unit. Service action codes available are:		
	Code	Name	Description
	00h	Register	Register a reservation key with the device server
	01h	Reserve	Create a persistent reservation using a reservation key
	02h	Release	Release a persistent reservation
	03h	Clear	Clear all reservation keys and all persistent reservations
	04h	Pre-empt	Pre-empt persistent reservations from another initiator
	05h	Pre-empt & Clear	Pre-empt persistent reservations from another initiator and clear the task set for the pre-empted initiator
	06 – 1Fh	Re-served	Reserved
	Refer to Table 5-114 for detailed descriptions of each of the service action codes.		
Scope	The value in this field indicates whether a persistent reservation applies to an entire logical unit, to a part of the logical unit (defined as an extent), or to an element. The values for the Scope field are:		
	Code	Name	Description
	0h	LU	Logical Unit. Persistent reservation applies to the full logical unit. The LU scope is therefore implemented by all device servers that implement PERSISTENT RESERVATION OUT.
	1h	Extent	Persistent reservation applies to the specified extent. When Extent is the scope, it indicates that the persistent reservation applies to the extent of the logical unit defined by the Scope-Specific Address and Extent Length fields in the PERSISTENT RESERVATION OUT command's parameter list. Note that an extent is defined only for devices defining contiguous logical block addresses. The Extent scope is optional for all device servers that implement PERSISTENT RESERVATION OUT.
	2h	Element	Persistent reservation applies to the specified element. When Element is the scope, it indicates that the persistent reservation applies to the element of the logical unit defined by the Scope-Specific Address field in the PERSISTENT RESERVATION OUT parameter list. An element is as defined by the SCSI-3 Medium Changer Commands (SMC) standard. The Element scope is optional for all device servers that implement PERSISTENT RESERVATION OUT.
	3h – Fh	Re-served	Reserved
Type	The value of the Type field specifies the characteristics of the persistent reservation being established for all data blocks within the extent or within the logical unit. Refer to Table 5-115 for the applicable Type codes and their meanings		
Parameter List Length	Fields contained in the PERSISTENT RESERVATION OUT parameter list specify the reservation keys and extent information required to perform a persistent reservation service action. The parameter list is 24 bytes in length; the Parameter List Length field contains 24 (18h) bytes.		

Table 5-114 provides detailed descriptions of each of the PERSISTENT RESERVATION OUT command's seven possible service actions (Service Action codes appear in bits 0 – 4 of Byte 1).

Table 5-114 *PERSISTENT RESERVATION OUT Command's Service Action Descriptions*

Code	Name	Description
00h	Register	<p>When the command executes a Register service action, it registers a reservation key with a device server without generating a reservation. The device server holds these reservation keys from each initiator that performs a PERSISTENT RESERVATION OUT command with a Register service action until the key is changed by a new PERSISTENT RESERVATION OUT command with Register service action from the same initiator, or until the initiator registration is removed by:</p> <ul style="list-style-type: none"> • Powering down the logical unit, if the last Activate Persist Through Power Loss (APTPL; see Figure 5-55 and Table 5-61) received by the device server was 0; • Performing a Clear service action; • Performing a Pre-empt service action; • Performing a Pre-empt and Clear service action; or • Performing a Register service action from the same initiator with the value of the service action reservation key set to 0. <p>When a reservation key has not yet been established or when the reservation key has been removed, a reservation key of 0 is used when the initiator performs a PERSISTENT RESERVATION OUT with the Register service action. When the reservation has been removed, no information is reported for the initiator in the Read Keys service action of the resulting PERSISTENT RESERVATION IN command.</p>
01h	Reserve	<p>A PERSISTENT RESERVATION OUT command with Reserve service action creates a persistent reservation with a specified Scope and Type. Persistent reservations are not superseded by a new persistent reservation from any initiator except by the execution of a PERSISTENT RESERVATION OUT command that specifies either a Pre-empt or Pre-empt and Clear service action. New persistent reservations that do not conflict with an existing persistent reservation execute normally. Persistent reservations of logical units or extents having the same Type value are permitted if no conflicting persistent reservations are held by another initiator. When these types of overlapping reservations are released, each of the extent reservations and logical unit reservations are removed with a separate Release service action. Multiple identical reservations from the same initiator are all released simultaneously via a single Release service action that matches the reservations.</p>

Code	Name	Description
02h	Release	<p>A PERSISTENT RESERVATION OUT command with Release service action removes a persistent reservation held by the same initiator.</p> <p>The fields associated with a Release service action match fields of the active persistent reservation. Sending of a PERSISTENT RESERVATION OUT command that specifies a Release service action when no persistent reservation exists from that initiator does not result in an error. Instead, the device server returns a GOOD message without altering any other reservation: the reservation key is not changed by the Release service action.</p> <p>The device server returns a CHECK CONDITION status for any PERSISTENT RESERVATION OUT command that specifies the release of a persistent reservation held by the requesting initiator that matches some but not all of the Scope, Type, Reservation Key, and extent values. The sense key is set to ILLEGAL REQUEST and additional sense data is set to INVALID RELEASE OF ACTIVE PERSISTENT RESERVATION. Attempts to release persistent reservations in which none of the Scope, Type, Reservation Key, and extent values match an existing persistent reservation held by the initiator making the request are not errors.</p> <p>An active persistent reservation may also be released by:</p> <ol style="list-style-type: none"> 1. Powering off. When the most recent APTPL value received by the device server is 0, a power-off performs a hard reset, clears all persistent reservations, and removes all registered reservation keys; or 2. Executing a PERSISTENT RESERVATION OUT command from another initiator with a persistent reserve service action of Pre-empt or Pre-empt and Clear. <p>Note that a Release service action should not be performed if any operations interlocked by the persistent reservation have not yet completed.</p>
03h	Clear	<p>A PERSISTENT RESERVATION OUT command with a successful Clear service action removes all persistent reservations for all initiators. All reservation keys are also removed. Any commands from any initiator that have been accepted by the device server as non-conflicting continue their normal executions.</p> <p>A UNIT ATTENTION condition is established for all initiators for the cleared logical unit. The sense key is set to UNIT ATTENTION; the additional sense data is set to RESERVATIONS PREEMPTED.</p> <p>Note that applications should not use the Clear action service except during recoveries associated with initiator or system reconfiguration, since data integrity may be compromised.</p>

Code	Name	Description
04h	Pre-empt	<p>A PERSISTENT RESERVATION OUT command with a successful Pre-empt service action removes all persistent reservations for all initiators that have been registered with the Service action Reservation key specified in the PERSISTENT RESERVATION OUT command's parameter list. A persistent reservation is also established for the pre-empting initiator. Any commands from any initiator that have been accepted by the device server as non-conflicting continue their normal executions. If a PERSISTENT RESERVATION OUT command is sent that specifies a Pre-empt service action and no persistent reservation exists for the initiator identified by the Service action Reservation key, it is not an error condition.</p> <p>A UNIT ATTENTION condition is established for the pre-empted initiators. The sense key is set to UNIT ATTENTION; the additional sense data is set to RESERVATIONS PRE-EMPTED. Commands that follow are subject to the persistent reservation restrictions set by the pre-empting initiator.</p> <p>The persistent reservation thus created by the pre-empting initiator is defined by the Scope and Type fields of the PERSISTENT RESERVATION OUT command and the corresponding fields of the command's parameter list.</p> <p>The registration keys for the pre-empted initiators are removed by the Pre-empt service action; the reservation key for an initiator that has performed a Pre-empt service action with its own Reservation key specified in the Service action Reservation key remains unchanged, although all other specified releasing actions and reservation actions are performed.</p> <p>Note that persistent reservations are not superseded by a new persistent reservation from any initiator except by the execution of a PERSISTENT RESERVATION OUT that specifies either the Pre-empt or the Pre-empt and Clear service actions. New persistent reservations that do not conflict with an existing persistent reservation execute normally. The persistent reservation of a logical unit or extents having the same Type value are permitted if no conflicting persistent reservations other than the reservations being pre-empted are held by another initiator.</p> <p>Note: Service Action (Code) 4 is not supported by Atlas 10K V.</p>

Code	Name	Description
05h	Pre-empt & Clear	<p>A PERSISTENT RESERVATION OUT command with a Pre-empt & Clear service action removes all persistent reservations for all initiators that have been registered with the Service action Reservation key specified in the PERSISTENT RESERVATION OUT command's parameter list. It also establishes a persistent reservation for the pre-empting initiator. Any commands from the initiators being pre-empted are terminated as if an ABORT TASK management function had been performed by the pre-empted initiator. If a PERSISTENT RESERVATION OUT command is sent that specifies a Pre-empt & Clear service action and no persistent reservation exists for the initiator identified by the Service action Reservation key, it is not an error condition. If the key is registered, however, the Clear portion of the action executes normally.</p> <p>A UNIT ATTENTION condition is established for the pre-empted initiators. The sense key is set to UNIT ATTENTION; the additional sense data is set to RESERVATIONS PRE-EMPTED. Commands that follow, and retries of commands that timed out because there were cleared are subject to the persistent reservation restrictions set by the pre-empting initiator.</p> <p>The persistent reservation thus created by the pre-empting initiator is defined by the Scope and Type fields of the PERSISTENT RESERVATION OUT command and the corresponding fields of the command's parameter list.</p> <p>The Pre-empt & Clear service action clears any ACA or CA condition associated with the initiator that is pre-empted and clears any tasks with an ACA attribute from that initiator. Any Asynchronous Event Reporting operations in progress that were initiated by the device server are unaffected by a Pre-empt and Clear service action.</p> <p>The reservation key for the other initiators pre-empted are removed by the Pre-empt & Clear service action. The reservation key for an initiator that has sent a Pre-empt & Clear action with its own reservation key specified in the service action's reservation key remains unchanged, although all other specified clearing actions, releasing actions, and reservation actions are performed.</p> <p>Persistent reservations are not superseded by a new persistent reservation from any initiator except via execution of a PERSISTENT RESERVATION OUT that specifies either the Pre-empt or Pre-empt & Clear service action. New persistent reservations not in conflict with an existing persistent reservation execute normally. The persistent reservation of a logical unit or extent that has the same Type value is permitted as long as no conflicting persistent reservations other than the reservations being pre-empted are permitted.</p>

Table 5-115 presents the definitions of the characters of the available “Type” values from the Type field of the PERSISTENT RESERVATION IN Read Reservations parameters.

Each of the codes provides handling instructions for READ operations, for WRITE operations, and for subsequent attempts to establish persistent reservations (referred to as “Additional Reservations Allowed” in the table).

Table 5-115 *Persistent Reservation Type Codes and Their Meanings*

Code	Name	Description
0h	READ Shared	<p>READS: Shared; any application client on any initiator may execute commands that perform transfers from the disk to the initiator.</p> <p>WRITES: Prohibited; any command from any initiator that performs a transfer from the initiator to the disk results in a reservation conflict.</p> <p>ADDITIONAL RESERVATIONS: Allowed; any initiator may reserve the logical unit, extents, or elements as long as the persistent reservations do not conflict with any reservations already known to the device server.</p>
1h	WRITE Exclusive	<p>READS: Shared; any application client on any initiator may execute commands that perform transfers from the disk to the initiator.</p> <p>WRITES: Exclusive; any command from any initiator other than the initiator that holds the persistent reservation that attempts a transfer to the disk results in a reservation conflict.</p> <p>ADDITIONAL RESERVATIONS: Allowed; any initiator may reserve the logical unit, extents, or elements as long as the persistent reservations do not conflict with any reservations already known to the device server.</p>
2h	READ Exclusive	<p>READS: Exclusive; any command from any initiator other than the initiator that holds the persistent reservation that attempts a transfer from the disk results in a reservation conflict.</p> <p>WRITES: Shared; any application client on any initiator may execute commands that perform transfers to the disk from the initiator.</p> <p>ADDITIONAL RESERVATIONS: Allowed; any initiator may reserve the logical unit, extents, or elements as long as the persistent reservations do not conflict with any reservations already known to the device server.</p>
3h	Exclusive Access	<p>READS: Exclusive; any command from any initiator other than the initiator holding the persistent reservation that attempts a transfer from the disk results in a reservation conflict.</p> <p>WRITES: Exclusive; any command from any initiator other than the initiator holding the persistent reservation that attempts a transfer to the disk results in a reservation conflict.</p> <p>ADDITIONAL RESERVATIONS: Restricted; any PERSISTENT RESERVATION OUT command with the Reserve service action from any initiator other than the initiator holding the persistent reservation results in a reservation conflict. The initiator that holds the persistent reservation can reserve the logical unit, extents, or elements as long as the persistent reservations do not conflict with any reservations already known to the device server.</p>

Code	Name	Description
4h	Shared Access	<p>READS: Shared; any application client on any initiator may execute commands that perform transfers from the disk to the initiator.</p> <p>WRITES: Shared; any application client on any initiator may execute commands that perform transfers to the disk from the initiator.</p> <p>ADDITIONAL RESERVATIONS: Restricted; any PERSISTENT RESERVATION OUT command with the Reserve service action from any initiator other than the initiator holding the persistent reservation results in a reservation conflict. The initiator that holds the persistent reservation can reserve the logical unit, extents, or elements as long as the persistent reservations do not conflict with any reservations already known to the device server.</p>
5h	WRITE Exclusive, Registrants Only	<p>READS: Shared; any application client on any initiator may execute commands that perform transfers from the disk to the initiator.</p> <p>WRITES: Exclusive; any command from an initiator that has not previously performed a Register service action with the device server that attempts a transfer to the disk results in a reservation conflict</p> <p>ADDITIONAL RESERVATIONS: Allowed; any initiator may reserve the logical unit, extents, or elements as long as the persistent reservations to not conflict with any reservations already known to the device server.</p>
6h	Exclusive Registrants Only	<p>READS: Exclusive; any command from an initiator that has not previously performed a Register service action with the device server that attempts a transfer from the disk results in a reservation conflict.</p> <p>WRITES: Exclusive; any command from an initiator that has not previously performed a Register service action with the device server that attempts a transfer to the disk results in a reservation conflict</p> <p>ADDITIONAL RESERVATIONS: Allowed; any initiator may reserve the logical unit, extents, or elements as long as the persistent reservations to not conflict with any reservations already known to the device server.</p>
7h - Fh	Reserved	N/A

The PERSISTENT RESERVATION OUT command requires a parameter list, illustrated in Table 5-116 and defined in Table 5-117. Each of the fields of the parameter list are sent for every PERSISTENT RESERVATION OUT command, even if the field is not required for the specific Service action and/or Scope values

Table 5-116 PERSISTENT RESERVATION OUT Parameter List-Data Format

Byte	Bit	7	6	5	4	3	2	1	0	
0 – 7	(MSB)	Reservation Key								(LSB)
8 – 15	(MSB)	Service Action Reservation Key								(LSB)
16 – 19	(MSB)	Scope-Specific Address								(LSB)
20		Reserved								APTPL
21		Reserved								

22 – 23	(MSB) Extent Length (LSB)
---------	---------------------------

Table 5-117 *PERSISTENT RESERVATION OUT Parameter List-Field Descriptions*

Data Field	Description
Reservation Key	This field contains an 8-byte token that is provided by the application client to the device server to identify which initiator is the source of the PERSISTENT RESERVATION OUT command. The device server verifies that the Reservation Key in the PERSISTENT RESERVATION OUT command matches the Reservation Key that is registered for the initiator from which the command is received. If there is no match, the device server returns a RESERVATION CONFLICT status. The Reservation Key of the initiator is valid for all Service action and Scope values.
Service Action Reservation Key	<p>This field contains information needed for 3 service actions: the Register service action, the Pre-empt service action, and the Pre-empt & Clear service action. The Service Action Reservation Key is ignored for all other service actions.</p> <p>For the Register service action, the Service Action Reservation Key field contains the new Reservation Key to be registered.</p> <p>For the Pre-empt and the Pre-empt & Clear service actions, the Service Action Reservation Key contains the reservation key of the persistent reservations that are being pre-empted. For the Pre-empt and the Pre-empt & Clear actions, any failure of the Service Action Reservation Key to match any registered keys results in the device server returning a RESERVATION CONFLICT status.</p>
Scope-Specific Address	<p>If the Scope is an Extent reservation, this field contains the Logical Block Address of the extent and the Extent Length field (bytes 22 – 23) contain the number of blocks in the extent.</p> <p>If the Scope signifies an Element reservation, the Scope-Specific Address field contains the Element Address, zero-filled in the most significant bytes to fit the field; the Extent Length is set to 0.</p> <p>If the Service action is Register or Clear, or if the Scope is a Logical Unit reservation, both the Scope-Specific Address and Extent Length fields are set to 0.</p>
APTPL	<p>Activate Persist Through Power Loss. This bit is valid only for Register service actions; it is ignored for all other types of service actions.</p> <p>Support for APTPL set to 1 is optional. If a device server that does not support APTPL receives a 1 in that bit in a Register service action, it returns a CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and additional sense data set to INVALID FIELD IN PARAMETER LIST.</p> <p>If the last valid APTPL bit value received by the device server is 0, the loss of power in the target releases any persistent reservations and removes all reservation keys. If the last valid APTPL bit value is 1, the logical unit retains all persistent reservations and all reservation keys for all initiators even if power is lost and later returned. The most recently received valid APTPL value from any initiator governs the logical unit's behavior in the event of a power loss.</p>

Table 5-118 illustrates which fields are set by the application client and interpreted by the device server for each Service and Scope value.

Table 5-118 *Device Server Interpretation of Service and Scope Value*

Service Action	Allowed Scope	Parameters		
		Type	Service Action Reservation Key	Element or Element Parameters
Register	Ignored	Ignored	Valid	Ignored
Reserve	LU	Valid	Ignored	Ignored
Reserve	Extent	Valid	Ignored	Extent Valid
Reserved	Element	Valid	Ignored	Ignored
Release	LU	Valid	Ignored	Ignored
Release	Extent	Valid	Ignored	Extent Valid
Release	Element	Valid	Ignored	Element Valid
Clear	Ignored	Ignored	Ignored	Ignored
Clear	Ignored	Ignored	Ignored	Ignored
Pre-empt	LU	Valid	Valid	Ignored
Pre-empt	Extent	Valid	Valid	Extent Valid
Pre-empt	Element	Valid	Valid	Element Valid
Pre-empt & Clear	LU	Valid	Valid	Ignored
Pre-empt & Clear	Extent	Valid	Valid	Extent Valid
Pre-empt & Clear	Element	Valid	Valid	Element Valid

5.18 READ (6) Command (08h)

The READ (6) command requests that data be transferred from the drive to the initiator. The amount of data transferred is specified by the transfer length field. The starting point is specified by the logical sector address. The data format and field descriptions are shown in Table 5-119 and Table 5-120 respectively.

Table 5-119 READ (6) Command-Data Format

Byte	Bit	7	6	5	4	3	2	1	0
0		Operation Code (08h)							
1		Reserved			Logical Block Address				
2 - 3		Logical Block Address (continued)							
4		Transfer Length							
5		Control							

Table 5-120 READ (6) Command-Field Descriptions

Field	Description
Logical Block Address	Specifies the logical block address where the read operation will begin.
Transfer Length	Specifies the number of contiguous logical blocks of data to transfer. A transfer length of 0 indicates that 256 logical blocks will be transferred. Any other value indicates the number of logical blocks that will be transferred.

5.19 READ (10) Command (28h)

The READ (10) command requests that data be transferred from the drive to the initiator. The amount of data transferred is specified by the transfer length field. The starting point is specified by the logical sector address. The data format and field descriptions are shown in Table 5-121 and Table 5-122 respectively.

Table 5-121 READ (10) Command-Data Format

Byte	Bit	7	6	5	4	3	2	1	0
0		Operation Code (28h)							
1		Reserved			DPO	FUA	Reserved		RelAdr
2 - 5		Logical Block Address							
6		Reserved							
7 - 8		Transfer Length							
9		Control							

Table 5-122 READ (10) Command-Field Descriptions

Field	Description
DPO	Disable Page Out. When the value is 0, the drive caches the data read, using its normal Logical Read Unit policies. When the value is 1, the drive assigns to logical blocks accessed by the READ (10) command the lowest priority for being fetched into or retained by the cache.
FUA	Force Unit Access. When the value is 0, the drive is allowed to satisfy the read with cache-resident data. A value of 1 indicates that the drive accesses the media in performing the command before returning GOOD status.
RelAdr	Relative Addressing. The drive does not support relative addressing. Must be 0
Logical Block Address	Specifies the logical sector address where the read operation will begin.
Transfer Length	Specifies the number of contiguous logical blocks of data to transfer. A transfer length of 0 indicates that no logical blocks are to be transferred. Any other value indicates the number of logical blocks to be transferred. The drive returns a CHECK CONDITION status with the correct sense data if the specified transfer extends beyond the logical sector space.

5.20 READ BUFFER Command (3Ch)

The READ BUFFER command is used in conjunction with the WRITE BUFFER command. It allows the initiator to gain access to the data storage areas within the drive's controller. In particular, these two commands are used to check the integrity of the buffer-initiator data path, independent of the media, typically for fault isolation in the event of a failure. In addition, these commands can be used to verify the caching policies of the drive controller or other kinds of diagnostic and verification work.

Both the READ BUFFER and the WRITE BUFFER commands have a buffer ID and a buffer offset data field defined in their Control Descriptor Block. In the Combined Header and Data Mode, if both are set to 0, reference is made to a single 512-byte buffer set aside in the drive controller's data memory (independent of cache). This allows an initiator to verify the data path from the controller's data memory to the initiator without affecting the drive's internal data caches. The data format and field descriptions are shown in Table 5-123 and Table 5-124 respectively.

Table 5-123 READ BUFFER Command-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (3Ch)							
1	Reserved				Mode			
2	Buffer ID							
3 - 5	Buffer Offset							
6 - 8	Allocation Length							
9	Control							

Table 5-124 READ BUFFER Command-Field Descriptions

Field	Description
Mode	Four modes are supported by the drive; Combined Header and Data, Data Only, Descriptor Only, and Read Data from Echo Buffer. The permissible values for Buffer ID, Buffer Offset, are as follows:
0000b	Combined Header and Data – In this mode, data to be transferred is preceded by a 4-byte header that contains the buffer capacity. The resulting descriptor and data are transferred to the initiator during the command's data-in buffer transfer. <i>Buffer ID</i> must be 0 <i>Buffer Offset</i> must be 0 Allocation length must be set to between 0 and 516 (4 bytes of header and up to 512 bytes of data).
0010b	Data Only– In this mode, the data-in buffer transfer contains data from the controller's memory. <i>Buffer ID</i> . When non-zero, the value specifies a specific track line in the controller's cache. If Buffer ID of zero is specified, the buffer referenced is the same as if a Buffer ID of one were specified. <i>Buffer Offset</i> . When non-zero, the value of Buffer Offset can range from 0 to one less than the number of sectors/track multiplied by 512. This value must be either a zero or a multiple of 512. <i>Allocation Length</i> Must be between zero and the maximum number of sectors/track multiplied by 512.
0011b	Descriptor Only When requesting a descriptor (a description of the buffer referenced in the Buffer ID field), the resulting data is the offset boundary requirements and the length of each buffer. <i>Buffer ID</i> When non-zero, the value specifies a specific track line within the controller's cache. If Buffer ID of 0 is specified, the buffer referenced is the same as if a Buffer ID of 1 were specified. <i>Buffer Offset</i> Must be zero <i>Allocation Length</i> Must be between 4 (header) and the maximum number of sectors/track multiplied by 512 + header size.
1010b	Read Data from Echo Buffer In this mode, the drive transfers data to the host from the echo buffer. The echo buffer will transfer the same data as when the WRITE BUFFER command with the mode field set to echo buffer was issued. <i>Buffer ID</i> Ignored. <i>Buffer Offset</i> Ignored.

5.21 READ CAPACITY Command (25h)

The READ CAPACITY command reports formatted capacity and track-boundary information. The data format and field descriptions are shown in Table 5-125 and Table 5-126 respectively.

Table 5-125 READ CAPACITY Command-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (25h)							
1	Reserved							RelAdr
2 - 5	Logical Block Address							
6 - 7	Reserved							
8	Reserved							PMI
9	Control							

Table 5-126 READ CAPACITY Command-Field Descriptions

Field	Description
RelAdr	Relative Addressing is not supported by the drive; the field must be 0.
Logical Block Address	Used in conjunction with the PMI bit. Must be 0 when the PMI bit is 0.
PMI	<p>Partial Media Indicator. When the value in the field is 0, the Returned Logical Block Address (Table 5-127) is that of the last logical sector of the logical unit.</p> <p>When the value is 1, the Returned Logical Block Address is that of the last sector on the same cylinder as the Logical Block Address supplied in the Command Descriptor Block's LBA field. That is, the Returned Logical Block Address is that of the Logical Block Address before a substantial delay in data transfer will be encountered).</p> <p>In either case, Block Length in Bytes (Table 5-127) is the unit's logical sector size, as determined by the Block Length field in the Mode parameter Block Descriptor.</p>

Table 5-127 shows the READ CAPACITY Returned Data - Data Format that will be sent during the data-in buffer transfer for the command. Refer to the PMI description in Table 5-126 for a description of these data fields.

Table 5-127 READ CAPACITY Returned Data-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0 - 3	Returned Logical Block Address							
4 - 7	Block Length in Bytes							

5.22 READ DEFECT DATA Command (10) (37h)

The READ DEFECT DATA (10) command directs the drive to report defect information about the drive's medium. Refer to the FORMAT UNIT command for information about defects. The data format and field descriptions are shown in Table 5-128 and Table 5-129 respectively. The data sent to the host in the command's data-in buffer transfer consists of a four-byte Defect List header (Table 5-132 and Table 5-133), followed by any Defect Descriptors (Table 5-130 and Table 5-131). The READ DEFECT DATA (10) Command Descriptor Block is shown in Table 5-136; the data fields are described in Table 5-129.

Table 5-128 READ DEFECT DATA (10) Command-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (37h)							
1	Reserved							
2	Reserved		Plist	Glist	Defect List Format			
3 - 6	Reserved							
7 - 8	(MSB) Allocation Length (LSB)							
9	Control							

Table 5-129 *READ DEFECT DATA (10) Command-Field Description*

Field	Description	
PList	When the value is 0, the returned data does not include the Primary Defect List.	
	When the value is 1, the returned data includes the Primary Defect List. The defects are returned in ascending order.	
GList	When the value is 0, the returned data does not include the Grown Defect List.	
	When the value is 1, the returned data includes the Grown Defect List. The defects are returned in order of occurrence.	
	When both PList and GList are set to 1, the Primary Defect List and the Grown Defect List are returned. The PList is returned followed by the GList (the lists are not merged). When both are set to 0, only the Defect List Header (Table 5-132) is reported to the initiator.	
Defect List Format	This field specifies the format in which the defect descriptors are to be returned. The supported formats are:	
	000b	Block Format (Table 5-130). ¹
	100b	Bytes from the Index (Figure 5-63). The data returned from this mode does not include defects outside normal user space.
	101b	Physical Sector Format (Table 5-131). Default format. The defects returned from this mode do not include defects outside normal user space.
Allocation Length	If the number of defect descriptors the SCSI device has assigned does not exceed the capability of the ALLOCATION LENGTH field size but contains a value that is insufficient to transfer all of the defect descriptors, the defect list length shall not be adjusted to reflect the truncation and the device server shall not create a CHECK CONDITION status. The initiator is responsible for comparing the defect list length and the allocation length to determine that a partial list was received. If the number of defect descriptors the SCSI device has assigned exceeds the capability of the ALLOCATION LENGTH field size, the drive shall transfer no data and return a CHECK CONDITION status and the sense key shall be set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CBD.	

¹ All in-line spared blocks do not have an associated LBA, and, therefore, will not be seen in Logical Block Format.

Table 5-130 shows the Block Format returned when the Defect List Format is 000b

Table 5-130 *Defect Descriptor-Block Format.*

Bit Byte	7	6	5	4	3	2	1	0
0 - 3	Defective Block Address							

Table 5-131 shows the Defective Bytes From Index Format when the Defect List Format is 100b or the Physical Sector Format when the Defect List Format is 101b.

Table 5-131 *Defect Descriptor-Bytes From Index Format, or Physical Sector Format*

Bit Byte	7	6	5	4	3	2	1	0
0 - 2	Cylinder Number of Defect							
3	Head Number of Defect							
4 - 7	Defective Bytes From Index / Defective Sector Number							

Table 5-132 *Defect List Header -Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0	Reserved							
1	Reserved			Plist	Glist	Defect List Format		
2 - 3	Defect List Length							

Table 5-133 *Defect List Header-Field Descriptions*

Field	Description
PList	Primary Defect List. When the value is 0, the returned data does not include the Primary Defect List. When the value is 1, the returned data includes the Primary Defect List.
GList	Grown Defect List. When the value is 0, the returned data does not include the Grown Defect List. When the value is 1, the returned data includes the Grown Defect List.
Defect List Format	This field indicates the format of the descriptors returned. If this differs from the format requested in the Command Descriptor Block, a recovered error is returned.
Defect List Length	This field indicates the actual number of bytes that follow the defect list header.

5.23 READ DEFECT DATA Command (12) (B7h)

The READ DEFECT DATA (12) command directs the drive to report defect information about the drive's medium. Refer to the FORMAT UNIT command for information about defects. The data sent to the host in the command's data-in buffer transfer consists of a four-byte Defect List header (Table 5-138 and Table 5-139), followed by any Defect Descriptors (Table 5-136 and Table 5-137). The READ DEFECT DATA (12) Command Descriptor Block is shown in Table 5-134; the data fields are described in Table 5-135.

Table 5-134 READ DEFECT DATA (12) Command-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (B7h)							
1	Reserved			Plist	Glist	Defect List Format		
2 - 5	Reserved							
6 - 9	(MSB) Allocation Length (LSB)							
10	Reserved							
11	Control							

Table 5-135 READ DEFECT DATA (12) Command-Field Description

Field	Description
PList	When the value is 0, the returned data does not include the Primary Defect List. When the value is 1, the returned data includes the Primary Defect List. The defects are returned in ascending order.
GList	When the value is 0, the returned data does not include the Grown Defect List. When the value is 1, the returned data includes the Grown Defect List. The defects are returned in order of occurrence. When both PList and GList are set to 1, the Primary Defect List and the Grown Defect List are returned. The Plist is returned followed by the Glist (the lists are not merged). When both are set to 0, only the Defect List Header (Table 5-138) is reported to the initiator.
Defect List Format	This field specifies the format in which the defect descriptors are to be returned. The supported formats are:
	000b Block Format (Table 5-136). ¹
	100b Bytes from the Index (Table 5-137). The data returned from this mode does not include defects outside normal user space.
	101b Physical Sector Format (Table 5-137). Default format. The defects returned from this mode do not include defects outside normal user space.
Allocation Length	If the number of defect descriptors the SCSI device has assigned does not exceed the capability of the ALLOCATION LENGTH field size but contains a value that is insufficient to transfer all of the defect descriptors, the defect list length shall not be adjusted to reflect the truncation and the device server shall not create a CHECK CONDITION status. The initiator is responsible for comparing the defect list length and the allocation length to determine that a partial list was received. If the number of defect descriptors the SCSI device has assigned exceeds the capability of the ALLOCATION LENGTH field size, the drive shall transfer no data and return a CHECK CONDITION status and the sense key shall be set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

¹ All in-line spared blocks do not have an associated LBA, and, therefore, will not be seen in Logical Block Format.

Table 5-136 shows the Block Format returned when the Defect List Format is 000b.

Table 5-136 *Defect Descriptor-Block Format*

Bit Byte	7	6	5	4	3	2	1	0
0 - 3	Defective Block Address							

Table 5-137 shows the Defective Bytes From Index Format when the Defect List Format is 100b or the Physical Sector Format when the Defect List Format is 101b.

Table 5-137 *Defect Descriptor-Bytes From Index Format, or Physical Sector Format*

Bit Byte	7	6	5	4	3	2	1	0
0 - 2	Cylinder Number of Defect							
3	Head Number of Defect							
4 - 7	Defective Bytes From Index / Defective Sector Number							

Table 5-138 *Defect List Header-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0	Reserved							
1	Reserved			Plist	Glist	Defect List Format		
2 - 3	Reserved							
4 - 7	Defect List Length							

Table 5-139 *Defect List Header-Field Descriptions*

Field	Description
PList	Primary Defect List. When the value is 0, the returned data does not include the Primary Defect List. When the value is 1, the returned data includes the Primary Defect List.
GList	Grown Defect List. When the value is 0, the returned data does not include the Grown Defect List. When the value is 1, the returned data includes the Grown Defect List.
Defect List Format	This field indicates the format of the descriptors returned. If this differs from the format requested in the Command Descriptor Block, a recovered error is returned.
Defect List Length	This field indicates the actual number of bytes that follow the defect list header.

5.24 READ LONG Command (3Eh)

The READ LONG command is used in conjunction with the WRITE LONG command to allow the initiator to issue raw reads and writes of physical data blocks, but address them as logical blocks with an assumed sector length of 512 bytes. These two commands are used to verify the correction capability of the drive's ECC and to create a known kind and location to verify the drive's recovery and bad sector handling policies. The data format and field descriptions are shown in Table 5-140 and Table 5-141 respectively.

Table 5-140 READ LONG Command Descriptor Block-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (3Eh)							
1	Reserved						Correct	RelAdr
2 - 5	Logical Block Address							
6	Reserved							
7 - 8	Byte Transfer Length (578)							
9	Control							

Table 5-141 READ LONG Command Descriptor Block-Field Descriptions

Field	Description
Correct	When the value in this field is 1, the drive corrects the data by the ECC before transferring it to the initiator. When the value is 0, the drive does not perform ECC correction on the data. NOTE: This bit, not the DCR bit of the Read Recovery page, governs the effect of error correction for the READ LONG command.
RelAdr	Relative addressing is not supported by the drive. This value must be 0.
Logical Block Address	Specifies the logical sector to read.
Byte Transfer Length	Must be 578. This value reflects the number of bytes spanned by the ECC, including the ECC itself, rounded up to the next word (this is easiest understood in bits: 4096 data bits + 1 force error bit + 15 Logical Block Address bits + 16 cross check bits + 2 pad bits + 480 ECC bits + 6 pad bits that are not written to media but are only needed for transfer = 4496 bits or 578 bytes). Any other value in this field results in a CHECK CONDITION status and the appropriate sense data.

Table 5-142 shows the format of the data that is returned during the READ LONG Command's data-in buffer transfer. The returned data always consists of 512 bytes of data, followed by a two LBA bytes, two EDC bytes, and 60 ECC-related bytes.

Table 5-142 READ LONG Command-Returned Data

Bit Byte	7	6	5	4	3	2	1	0
0 - 511	Data (512 Bytes)							
512 - 513	Logical Block Address Bits + Force Error Bit							
514 - 515	EDC (8 Bits)							
516 - 577	60 ECC Bits + 8 Pad Bits							

5.25 READ SKIP MASK Command (E8h)

The READ SKIP MASK command allows the initiator to condition the function of the READ commands to allow only selected data to be transferred over the SCSI interface. A READ SKIP MASK command precedes a READ (10) command; the READ (10) command must be linked to the READ SKIP MASK command or the skip mask is lost and a CHECK CONDITION status is returned.

Table 5-143 *READ SKIP MASK Command-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (E8h)							
1	(MSB)	Reserved						(LSB)
2 - 5	Logical Block Address							
6	Skip Mask Length							
7 - 8	(MSB)	Transfer Length						(LSB)
9	VU	Reserved					Flag	Link

Table 5-144 *READ SKIP MASK Command-Field Description*

Field	Description
Logical Block Address	This field specifies the first logical block of the corresponding READ (10) COMMAND. This allows a SEEK to begin immediately. The LBA of the following READ (10) command must be the same as the LBA of the READ SKIP MASK command.
Skip Mask Length	This field indicated the length in bytes of the skip mask. The mask is up to 64 bytes, giving a data address capability of 256K. The mask consists of a sequence of bits where a "1" bit is for the blocks of data to be transferred and a "0" for the blocks of data that are to be skipped. The first "1" bit of the mask corresponds to the first LBA to be transferred. Any unused bits at the end of the last byte in the mask must be set to 0. The mask is transferred in the data-out buffer transfer of the skip mask command.
Transfer Length	This field must be equal to the transfer length in the linked READ (10) command. The transfer length field must be equal to the number "1" bits in the skip mask. If a READ (10) is used with 0 as the transfer length, a transfer length of 256 data blocks is implied.

5.26 REASSIGN BLOCKS Command (07h)

The REASSIGN BLOCKS Command reassigns one or more logical blocks to a spare location(s) set aside for this purpose and records the reassignments in the Grown Defect List.

No command-specific Control Descriptor Block fields affect the processing performed for this command.

The list of Logical Block Addresses to reassign is supplied in the REASSIGN BLOCKS Defect List passed in the command's data-out buffer transfer. The Defect List consists of a Defect List Header (Table 5-146) followed by zero or more Defect Descriptors (Table 5-148). Table 5-147 contains the field description for the Defect List Header. If multiple Logical Block Addresses are specified, they are listed in ascending order.

The REASSIGN BLOCKS Command Descriptor Block is shown in Table 5-145.

Table 5-145 REASSIGN BLOCKS Command Descriptor Block-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (07h)							
1 - 4	Reserved							
5	Control							

Table 5-146 REASSIGN BLOCKS Defect List Header-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0 - 1	Reserved							
2 - 3	Defect List Length							

Table 5-147 REASSIGN BLOCKS Defect List Header -Field Description

Field	Description
Defect List Length	The contents of this field indirectly indicates the number of defective Logical Block Addresses passed in the Defect List. The value must be four times the number of Defect Descriptors supplied. Zero (0) is a valid value and makes the command a NOP.

Table 5-148 REASSIGN BLOCKS Defect Descriptor-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0 - 3	Defect Logical Block Address							

The drive returns a CHECK CONDITION status if it cannot complete the command. If this happens because the drive ran out of spare sectors, the drive sets the sense key to MEDIUM ERROR and the additional sense code to NO DEFECT SPARE LOCATION AVAILABLE. The sense key and additional sense key are set appropriately for any other cause. The Logical Block Address of the first defect descriptor not reassigned is returned in the command-specific information field of the

sense data. If all defects were reassigned, but some other error occurred, then FFFFFFFFh is posted in the command-specific field.

5.27 RECEIVE DIAGNOSTIC RESULTS Command (1Ch)

The RECEIVE DIAGNOSTIC RESULTS command fetches the results of the last SEND DIAGNOSTIC command sent to the drive. The drives support the diagnostic pages listed in Table 5-149.

Table 5-149 *Diagnostic Pages Supported by The Drives*

Page Code	Description	Size In Bytes
00h	Supported Diagnostic Pages	6
40h	Translate Address Page (SEND DIAG)	14
40h	Translate Address Page (RECEIVE DIAG)	14 or more

The RECEIVE DIAGNOSTIC RESULTS Command Descriptor Block is shown in Table 5-150; the data fields are described in Table 5-151.

To receive one or both of the two diagnostic pages that are available, the initiator first sends the desired page code(s) to the drive using a SEND DIAGNOSTICS command. The SEND DIAGNOSTICS command tells the drive which page it should return in response to the next RECEIVE DIAGNOSTICS RESULTS command. The page specified in the most recent SEND DIAGNOSTICS command is always the page returned in response to the RECEIVE DIAGNOSTICS RESULTS command.

Table 5-150 *RECEIVE DIAGNOSTIC RESULTS Command Descriptor Block-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (1Ch)							
1	Reserved							PCV
2	Page Code Value							
3 - 4	Allocation Length							
5	Reserved			Control				

Table 5-151 *RECEIVE DIAGNOSTIC RESULTS Command Descriptor Block-Field Descriptions*

Field	Description
Allocation Length	Specifies the number of bytes of diagnostic-page results the drive is allowed to send to the initiator. A CHECK CONDITION is not returned if this value truncates the information available.

5.27.1 Supported Diagnostics Pages Page (00h)

Table 5-152 shows the format of the Supported Diagnostics Pages page.

Table 5-152 *Supported Diagnostics Pages Page-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0	Page Code (00h)							
1	Reserved							
2 - 3	Page Length (0002h)							
4	Supported Diagnostics Pages Page (00h)							
5	Translate Address Page (40h)							

To obtain the Supported Diagnostics Pages list, the initiator must:

1. Send a Supported Diagnostics Page to the drive with a SEND DIAGNOSTICS command (the Page Length in the Supported Diagnostics page must be 0).
2. Send a RECEIVE DIAGNOSTIC RESULTS command to the drive. The listing of Supported Diagnostic Pages is returned by the drive.

5.27.2 Translate Address Page (40h)

The Translate Address Page allows the initiator to translate a logical sector address, physical address, or physical bytes for an index address into any one of the other formats. The address to be translated is passed to the drive with the SEND DIAGNOSTIC command, and the results are returned to the initiator by the RECEIVE DIAGNOSTIC RESULTS command.

Table 5-153 shows the format of the Translate Address Page that is returned in response to a RECEIVE DIAGNOSTIC RESULTS command, and Table 5-154 contains a description of the data fields. Multiple translated addresses are returned if the logical sector size is greater than 512 bytes (more than one physical data sector).

Table 5-153 *Translate Address Page-Data Format*

Byte	Bit	7	6	5	4	3	2	1	0
0	Operation Code (40h)								
1	Reserved								
2 - 3	Page Length ($2 + (8 \times n)$) (where n = number of sectors / LBA)								
4	Reserved						Supplied Format = 000b (LBA)		
5	RAREA	ALT-SEC	ALT-TRK	Reserved			Translate Format = 101b (Physical Sector)		
6 - 8	Translated Address (Cylinder Number) ¹								
9	Translated Address (Head Number) ¹								
10 - 13	Translated Address (Sector Number or Bytes From Index) ¹								
y - y + 2	Translated Address (Cylinder Number) Where $y = ((n-1) \times 8) + 6$								
y + 3	Translated Address (Head Number)								
y + 4 - y + 7	Translated Address (Sector Number or Bytes From Index)								

Note: 1. Bytes 6 through 13: One set of these bytes is required for each sector per Logical Block Address.

Table 5-154 *Translate Address Page-Field Descriptions*

Field	Description
Supplied Format	The Supplied Format field specifies the format of the following Translate Address field. Must be 000b, which specifies the LBA format.
RAREA	Reserved Area is not used. <i>Must be 0.</i>
ALTSEC	An Alternate Sector bit of 1 indicates that the translated address from the SEND DIAGNOSTIC command is physically located in an alternate sector on the disk. When ALTSEC is 0, it indicates either that no part of the translated address is located in an alternate sector of the medium or that the drive is unable to determine this information.
ALTTRK	An Alternate Track bit of 1 indicates either that part or all of the translated address is located on an alternate track of the medium or, the drive cannot determine if all or part of the translated address is located on an alternate track. When ALTTRK is 0, it indicates that no part of the translated address is located on an alternate track of the medium.
Translate Format	The Translate Format field specifies the format into which the initiator wants the address translated. The value must be 101b, specifying the Physical Sector format.

5.28 RELEASE (6) Command (17h)

The RELEASE and the RESERVE commands are used for contention resolution in multiple-initiator systems. The RELEASE command is used to release a previously reserved logical unit. The drive will not return an error if the initiator attempts to release a reservation that is not currently valid. The data format and field descriptions are shown in Table 5-155 and Table 5-156 respectively.

Table 5-155 RELEASE (6) Command Descriptor Block-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (17h)							
1	Reserved			3rdPty	Third Party Device ID			Extent
2	Reservation Identification							
3 - 4	Reserved							
5	Control							

Table 5-156 RELEASE (6) Command -Field Descriptions

Field	Description
3rdPty	Third Party Release allows an initiator to release a logical unit that was previously reserved. If the value in this field is 0, third party release is not requested. When the value is 1, the drive releases the specified logical unit, but only if the reservation was made by the initiator using a third-party reservation that requested the release for the same SCSI device as specified in the Third Party Device ID field.
Third Party Device ID	Required if the 3rdPty bit is 1. This field specifies the SCSI ID (0 through 7 only) of the initiator whose third party reservation is being released. (This field must be set if the initiator of the original third party RESERVE is the source of the RELEASE.)
Extent	The drive supports reservations only on entire logical units. <i>The value must be 0.</i>
Reservation Identification	Any value in this field is ignored by the drive.

Note: Assuming that the RELEASE Command Descriptor Block is valid, the drive always returns a GOOD status for this command. An actual release only happens if the initiator has the unit reserved for itself or a third-party initiator.

5.29 RELEASE (10) Command (57h)

The RELEASE and the RESERVE commands are used for contention resolution in multiple-initiator systems. The RELEASE (10) command is used to release a previously reserved logical unit. The drive will not return an error if the initiator attempts to release a reservation that is not currently valid. The data format and field descriptions are shown in Table 5-157 and Table 5-158 respectively.

Table 5-157 RELEASE (10) Command Descriptor Block-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (57h)							
1	Reserved			3rdPty	Reserved		LongID	Extent
2	Reservation Identification							
3	Third Party Device ID							
4 - 6	Reserved							
7 - 8	(MSB) Parameter List Length (LSB)							
9	Control							

Table 5-158 RELEASE (10) Command — Field Descriptions

Field	Description
3rdPty	Third Party Release allows an initiator to release a logical unit that was previously reserved. If the value in this field is 0, third party release is not requested. If 3rdPty = 1, then the device server will release the specified logical unit, but only if the initiator ID, 3rdPty bit, and third party device ID are identical when compared to the RESERVE command that established the reservation.
LongID	If the Long ID bit is set to 1, the Parameter List Length is 8 and the eight bytes of the parameter list carry the device ID of the third party device; the contents of the Third Party Device ID in the CDB (byte 3) are ignored.
Third Party Device ID	If the Third Party Device ID value that is connected with the reservation release is smaller than 255, the LongID bit may be 0 and the ID value sent in the CDB. If LongID bit = 0, the Parameter List Length field also = 0. If the Third Party Device ID value is greater than 255, LongID = 1. Device servers that support device IDs greater than 255 will accept commands with LongID = 1; device servers that do not support IDs greater than 255 may reject commands with LongID set = 1. Device ID formats are protocol-specific.
Parameter List Length	The contents of this field specify the length, in bytes, of the parameter list that will be transferred from the initiator to the target.
Extent	The drive supports reservations only on entire logical units. <i>The value must be 0.</i>
Reservation Identification	Any value in this field is ignored by the drive.

Note: Assuming that the RELEASE Command Descriptor Block is valid, the drive always returns a GOOD status for this command. An actual release only happens if the initiator has the unit reserved for itself or a third-party initiator.

5.30 REPORT DEVICE IDENTIFIER Command (A3h)

The REPORT DEVICE IDENTIFIER command requests that the device server send device identification to the application client. The data format and field descriptions are shown in Table 5-159 and Table 5-160 respectively.

Table 5-159 REPORT DEVICE IDENTIFIER Command Descriptor Block-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (A3h)							
1	Reserved			Service Action (05h)				
2 – 5	Reserved							
6 – 9	(MSB) Allocation Length (LSB)							
10	Reserved							
11	Control							

Table 5-160 REPORT DEVICE IDENTIFIER Command Descriptor Block-Field Descriptions

Field	Description
Service Action	Must = 05h. Any other value forces Check Condition, Illegal Request.
Allocation Length	If the length is not sufficient to contain all of the parameter data, the first portion of the data is returned, and this is not considered to be an error. The actual length of the parameter data is available in the Identifier Length field in the parameter data. If the remainder of the parameter data is required, the application client must send a new REPORT DEVICE IDENTIFIER command with an Allocation Length field large enough to contain all of the data.

The REPORT DEVICE IDENTIFIER command's parameter list contains a 4-byte field that provides the length, in bytes, of the parameter list and the logical unit's identifier. The data format and field descriptions are shown in Table 5-161 and Table 5-162 respectively.

Table 5-161 REPORT DEVICE IDENTIFIER Parameter List-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0 – 3	(MSB) Identifier Length (n – 4)							(LSB)
4 – n	Identifier							

Table 5-162 REPORT DEVICE IDENTIFIER Parameter List-Field Descriptions

Field	Description
Identifier Length	This field specifies the length, in bytes, of the Identifier field. If the Allocation Length field in the CDB is insufficient in size to transfer all of the Identifier, the length is not adjusted to reflect any truncation of information. The Identifier is initially set = 0, and is changed only by a successful SET DEVICE IDENTIFIER command.

Identifier	The value reported in this field is a vendor-specific value and is the last value written by a successful SET DEVICE IDENTIFIER command. The value of the Identifier is changed only by a successful SET DEVICE IDENTIFIER command. The Identifier value persists through device resets, power cycles, and formatting of media.
------------	---

5.31 REPORT LUNS Command (A0h)

The REPORT LUNS command requests that the peripheral device logical unit numbers of known logical units in the target disk be sent to the applications client.

The command only returns information about the logical units to which commands may be sent. The data format and field descriptions are shown in Table 5-163 and Table 5-164 respectively.

Table 5-163 *REPORT LUNS Command Descriptor Block-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (A0h)							
1 - 5	Reserved							
6 - 9	(MSB) Allocation Length (LSB)							
10	Reserved							
11	Control							

Table 5-164 *REPORT LUNS Command Descriptor Block-Field Description*

Field	Description
Allocation Length	If the Allocation Length is not sufficient to contain the logical unit number values for all configured logical units, the device server still reports as many logical number values as will fit in the Allocation Length. The format of the report of configured logical units is shown in Table 5-165.

Table 5-165 *LUN Reporting Parameter List -Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0 - 3	(MSB) LUN List Length (n - 7) (LSB)							
4 - 7	Reserved							
8 - 15	(MSB) LUN (first LUN) (LSB)							
n-7 - n	LUN (last LUN, if more than one)							

The LUN List Length field contains the length in bytes of the LUN list that can be transferred. The LUN list length equals the number of logical unit numbers reported multiplied by eight. If the allocation length in the CDB is too small to allow transfer of information about all of the logical units configured, the LUN list length value is not adjusted to reflect the truncation. The data format for the LUN Reporting Parameter List is described in Table 5-165.

5.32 REQUEST SENSE Command (03h)

The REQUEST SENSE command causes the drive to transfer detailed sense data to the initiator. The drive maintains sense data on an individual initiator basis. The data format and field descriptions are shown in Table 5-166 and Table 5-167 respectively.

Table 5-166 REQUEST SENSE Command Descriptor Block-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (03h)							
1 - 3	Reserved							
4	Allocation Length							
5	Control							

Table 5-167 REQUEST SENSE Command Descriptor Block-Field Description

Field	Description
Allocation Length	The drives can return a maximum of 18 bytes of sense data. The initiator should set this field to the maximum value to receive all the sense data.

5.32.1 Sense Data Availability

Sense data is available if the initiator is responsible for an in-progress FORMAT UNIT command (Format Progress Indication is available in the Format Progress Indication Bytes).

The drive returns a sense key of NO SENSE and an additional sense code of NO ADDITIONAL SENSE INFORMATION if it has no sense data available for the initiator.

For Parallel SCSI disk drives, sense data is available if the previous command from the initiator terminated with a CHECK CONDITION status, or if the previous command from the initiator ended with an unexpected BUS FREE error. REQUEST Clearing Sense Data

Sense data is cleared from the drive's memory by one of the following conditions:

- After being returned from a REQUEST SENSE command.
- As soon as any other command is received from the initiator except INQUIRY.

5.32.2 Status Reporting

The drive returns CHECK CONDITION status for a REQUEST SENSE command only to report errors specific to the command when:

- A non-zero reserved bit is detected in the Command Descriptor Block.
- A drive malfunction prevents return of the sense data.

For Parallel SCSI disk drives, the drive also returns CHECK CONDITION status for a REQUEST SENSE command to report errors specific to the command when an unrecovered parity error is detected on the data bus.

Note that any of the previous conditions overwrites the current sense data with the sense data describing the error encountered by the REQUEST SENSE command (that is, for the condition stated in the first bulleted entry, the new sense data reflects that a non-zero reserved bit was detected in the Command Descriptor Block).

The drive returns a maximum of 18 bytes of sense data. The initiator should set the Allocation Length field in the Command Descriptor Block to 18 to receive all sense data. If the Allocation Length is set to 0, four bytes of data will be returned (this option maintains compatibility with SCSI-1 systems).

5.32.3 Sense Data Format for Error Code 70h and Error Code 71h

The sense data format for error code 70h (current error) and error code 71h (deferred error) is shown in Table 5-168; the data fields are defined in Table 5-169. The Supported Sense Keys are described in Table 5-170 and the field contents in Table 5-171. The Supported Additional Sense Codes and Sense Code Qualifiers are described in Table 5-172 and the Sense-Key Specific Field Contents in Table 5-173.

Table 5-168 Sense Data Format for Error Code 70h or 71h-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0	Valid	Error Code (70h or 71h)						
1	Segment Number = 00h							
2	FileMrk = 0	EOM = 0	ILI	Rsvd	Sense Key			
3 - 6	Information							
7	Additional Sense Length = 0Ah							
8 - 11	Command-Specific Information							
12	Additional Sense Code							
13	Additional Sense Code Qualifier							
14	Field Replaceable Unit Code							
15	SKSV	Sense-Key Specific						
16 - 17	Sense-Key Specific (continued)							

Table 5-169 *Sense Data Fields (Error Code 70h)-Field Descriptions*

Field	Definition	
Valid	The Valid bit, when set to 1, indicates that the Information field contains valid information, as defined below. When set to 0, it indicates that the contents of the Information field are not as defined below.	
Error Code	Either 70h, indicating that current errors are to be reported, or 71h, indicating deferred errors are to be reported.	
Segment Number	This feature is not supported and will return a zero value.	
FileMrk	The File Mark field is reserved for sequential access devices. This field will return a zero value.	
EOM	The End-of-Medium field is reserved for sequential access and printer devices. This field will return a zero value.	
ILI	Incorrect Length Indicator. When this bit is set to 1, it indicates that the requested logical sector length did not match the logical sector length of the data on the media. This field is only set for incorrect byte transfer lengths on a READ LONG and WRITE LONG command descriptor block.	
Sense Key	<p>The sense key, additional sense code, and additional sense code qualifier provide a hierarchy of information. The sense key provides the highest level of information about error and exception conditions. The additional sense code and qualifier provided more detailed information.</p> <p>The sense keys supported by the drive are described in Table 5-173.</p>	
Information	<p>The Information field contains the following device-type or command-specific information:</p> <ol style="list-style-type: none"> 1) The unsigned logical sector address associated with the sense key. 2) The difference between the requested length and the actual length in either bytes or sectors, as determined by the command. This difference is called the residue. 3) A description of each of the drive's supported commands definition for this field is contained in Table 5-171. 	
Additional Sense Length	The Additional Sense Length field indicates the number of additional sense bytes that can follow. If the Allocation Block Length specified in the Command Descriptor Block is too small to transfer all of the additional sense bytes, the Additional Sense Length byte value is not adjusted to reflect the truncation.	
Command-Specific Information	The Command-Specific Information field is valid only for Sense Data for a failed REASSIGN BLOCKS command. It contains either the Logical Block Address of the first defect descriptor not re-assigned or, if all the defect descriptors were reassigned, FFFFFFFh.	
Additional Sense Code	The Additional Sense Code field indicates further information related to the error or exception condition reported in the sense key. If the drive does not have further related information, the additional sense key is set to NO ADDITIONAL SENSE INFORMATION. See Table 5-172.	
Additional Sense Code Qualifier	The Additional Sense Code Qualifier field indicates detailed information related to the additional sense code. If the drive does not have detailed information, the field is set to 0.	
Field Replaceable Unit	The Field Replaceable Unit Code specifies a unit that has failed:	
	Code	FRU
	1	Host
	2	Hard Disk Assembly (HDA)
	3	Printed Circuit Board Assembly (PCBA)
SKSV	Sense-Key Specific Valid bit. This bit is 1 for all sense keys listed in Table 5-170 (exception: NOT READY and UNIT ATTENTION/RPL Status Change). The SKSV bit and Sense-Key Specific bytes may be 0 when the sense key is NOT READY.	
Sense-Key Specific	The contents of the Sense-Key Specific field are unique to each sense key. Table 5-171 summarizes the sense keys for which the drive supplies sense-key specific information. (See SKSV.)	

Table 5-170 *Supported Sense Keys*

Sense Key	Code	Description
NO SENSE	0h	There is no specific sense key information to be reported for the designated logical unit. This is the case for a successful command or a command that received CHECK CONDITION or COMMAND TERMINATED status because one of the following bits is set to 1: File Mark, EOM, or ILL.
RECOVERED ERROR	1h	The last command completed successfully with some recovery action performed by the drive. Details may be contained in the additional sense bytes and the Information field. When multiple recovered errors occur during one command, the choice of which error to report (first, last, most severe, etc.) is device specific.
NOT READY	2h	The logical unit cannot be accessed. Operator intervention may be required to correct this condition.
MEDIUM ERROR	3h	The command terminated with a non-recovered error condition that was probably caused by a flaw in the medium or an error in the recorded data. This sense key can also be returned if the drive is unable to distinguish between a flaw in the medium and a specific hardware failure (Sense Key
HARDWARE ERROR	4h	The drive detected a non-recoverable hardware failure while performing the command or during a self-test. The error type could be controller failure, device failure, etc.
ILLEGAL REQUEST	5h	The drive detected an illegal parameter in the Command Descriptor Block or in the additional parameters supplied as data for some commands. If the drive detects an invalid parameter in the Command Descriptor Block, it terminates the program without altering the medium. If the drive detects an invalid parameter in the additional parameters supplied as data, the drive may have already altered the medium. This sense key may also indicate that an invalid IDENTITY message was received.
UNIT ATTENTION	6h	The drive generates this sense key whenever the drive is reset by a TARGET RESET message, a hard reset condition, or by a power-on reset. Other conditions that generate a UNIT ATTENTION include: <ul style="list-style-type: none">• The mode parameters in effect for this initiator have been changed by another initiator.• The microcode has been changed.• Tagged commands queued for this initiator were cleared by another initiator.• INQUIRY data has been changed.• The mode parameters in effect for the initiator have been restored from non-volatile memory.• Other events that require the attention of the initiator.
DATA PROTECT	7h	A Read or Write command was attempted on a sector that is protected for this operation. The requested operation was not performed.
VENDOR SPECIFIC	9h	This sense key is used to report specific Maxtor conditions.
ABORTED COMMAND	Bh	The drive aborted the command. The initiator may be able to reissue the command successfully.
MISCOMPARE	Eh	The source data did not match the data read from the medium (i.e., a mismatch).

Table 5-171 Sense Key Information Field Contents

Command	Sense Key	Field Contents
FORMAT UNIT	RECOVERED ERROR	The address of the first bad sector encountered during the format operation. This sector, and all other bad sectors found during a format operation, are automatically reallocated by the drive.
	ILLEGAL REQUEST	The first illegal defect descriptor encountered in the Defect Descriptor List.
READ(6)(10)	RECOVERED/MEDIUM ERROR	The address of the first bad sector found during the read (possible only when PER=1, DTE=0, and DCR=0.). For a RECOVERED ERROR, the last Logical Block Address in the error is returned; for a MEDIUM ERROR the first Logical Block Address is returned.
READ LONG	RECOVERED/MEDIUM ERROR	READ(6)(10) error recovery qualifiers apply. The information field echoes this command's Logical Block Address field.
	ILLEGAL REQUEST	The transfer length residue (requested, minus correct length) if the transfer length is wrong. <i>Note: ILI will be set.</i>
REASSIGN BLOCKS	HARDWARE ERROR	The address of the first logical sector is not reassigned (the drive ran out of spare sectors).
	ILLEGAL REQUEST	The first illegal defect descriptor encountered in the Defect Descriptor List.
SEEK(6)(10)	ILLEGAL REQUEST	The first invalid sector address detected by the command.
VERIFY	RECOVERED/MEDIUM ERROR	WRITE(6)(10) error recovery qualifiers apply. Contains the Logical Block Address that may have failed the write or verify (compare) operation. For a RECOVERED ERROR, the last Logical Block Address in the error is returned; for a MEDIUM ERROR the first Logical Block Address is returned.
WRITE(6)(10)	RECOVERED/MEDIUM ERROR	The address of the first bad sector found during the write. (Possible only when PER=1, DTE=0, and DCR=0.) For a RECOVERED ERROR, the last Logical Block Address in the error is returned; for a MEDIUM ERROR the first Logical Block Address is returned.
WRITE AND VERIFY	RECOVERED/MEDIUM ERROR	WRITE(6)(10) error recovery qualifiers apply. This field contains the Logical Block Address that may have failed the write or verify (compare) operation. For a RECOVERED ERROR, the last Logical Block Address in the error is returned; for a MEDIUM ERROR the first Logical Block Address is returned.
WRITE LONG	RECOVERED/MEDIUM ERROR	Write (10) error recovery qualifiers apply. The information field echoes the WRITE LONG command's Logical Block Address.
	ILLEGAL REQUEST	The transfer length residue (requested minus correct length) if the transfer length is wrong. <i>Note: ILI will be set.</i>

Table 5-172 Supported Additional Sense Codes and Sense Code Qualifiers (in Hex)

Sense Key	Sense Code	Sense Qualifier	Meaning
0x00 No Sense	0x00	0x00	No additional sense information
	0x0B	0x01	SMART overtemperature trip
	0x5D	0x00	Smart triggered
	0x5D	0x01	Smart triggered due to Read TA Factor
	0x5D	0x02	Smart triggered due to Write Err Factor
	0x5D	0x5B	Smart triggered due to spin retry
	0x5D	0xFF	Smart test trigger
	0x5E	0x00	Power State: ACTIVE
	0x5E	0x01	Power State: IDLE
	0x5E	0x02	Power State: STANDBY
	0x5E	0x03	Power State: SLEEP
0x01 Rec Error	0x00	0x80	H-Grown: Backout prev Grown
	0x02	0x00	No seek complete: recovered
	0x03	0x00	Write fault: recovered
	0x03	0x01	Pre-Amp write unsafe: recovered
	0x09	0x00	Track following error: recovered
	0x0B	0x01	SMART overtemperature trip
	0x0C	0x00	Recovered write error
	0x0C	0x01	Write error - recovered with auto reallocation
	0x15	0x02	Positioning error detected by read: recovered
	0x15	0x80	Positioning error detected by read on write: recovered
	0x15	0x81	Servo Write Not OK assertion: recovered
	0x15	0x82	Servo Read Not OK assertion: recovered
	0x15	0x83	Spoke address error: recovered
	0x16	0x00	Data synchronization mark error: recovered
	0x16	0x83	Spoke synchronization mark error: recovered
	0x17	0x00	Recovered data with no ECC applied
	0x17	0x01	Recovered data with retries
	0x17	0x86	EC REC SVO MISSED STM
	0x17	0x91	TA Err-1st. - recovered
	0x17	0x92	TA Err-2nd. - recovered
	0x17	0x93	Bad Sync-1st. - recovered
	0x17	0x94	Bad Sync-2nd. - recovered
	0x17	0x95	Offtrack Prediction Err-1st. - recovered
	0x17	0x96	Offtrack Prediction Err-2nd. - recovered
	0x17	0x97	Head Spoke Num Miscompare-1 - recovered
	0x17	0x98	Head Spoke Num Miscompare-2 - recovered
	0x17	0x99	Timeout waiting burst/Spoke-1 - recovered
	0x17	0x9A	Timeout waiting burst/Spoke-2 - recovered

Sense Key	Sense Code	Sense Qualifier	Meaning
0x01 Rec Error	0x17	0xA0	Possible Bump-1st - recovered
	0x17	0xA1	Bump Err-2nd - recovered
	0x17	0xA2	Defective Wdg-1st - recovered
	0x17	0xA3	Defective Wdg-2nd - recovered
	0x17	0xA4	Map'd Wdg-1st - recovered
	0x17	0xA5	Map'd Wdg-2nd - recovered
	0x17	0xA6	Id Err-1st - recovered
	0x17	0xA7	Id Err-2nd - recovered
	0x17	0xA8	Unknown Wdg Err-1st - recovered.
	0x17	0xA9	Unknown Wdg Err-2nd - recovered
	0x17	0xAA	Shock detected - recovered
	0x17	0xAB	Shock Detected Err-2nd. - recovered
	0x17	0xAF	Recovered data with ECC on-the-fly
	0x17	0xB0	Confined sync timeout - recovered
	0x17	0xB2	Context Acquire Timeout - recovered
	0x17	0xB3	Context Acquire Data Addr Valid T/O - recovered
	0x17	0xB5	False SSMF - recovered
	0x18	0x00	Recovered data with ECC applied
	0x18	0x02	Recovered data - data auto-reallocated
	0x18	0x05	Recovered data - recommend reassignment (TA err)
	0x18	0x85	Recovered data - thermal asperity error
	0x1C	0x00	Requested Defect List Not Found
	0x40	0x85	EDC check failed - recovered
	0x40	0x86	Controller FIFO over/underflow error: recovered
	0x40	0x8C	LBA mismatch error: recovered.
	0x41	0x86	Required HW process did not complete: recovered
	0x42	0x00	Power On or Self Test failed
	0x5D	0x00	Smart triggered
	0x5D	0x01	Smart triggered due to Write TA Factor
	0x5D	0x02	Smart triggered due to Write Err Factor
	0x5D	0x5B	Smart triggered due to spin retry
	0x5D	0xFF	Smart test tripped
	0x80	0x10	Block Slow PLL Error: recovered.
	0x80	0x11	Split Slow PLL Error: recovered.
	0x80	0x12	Block state machine disabled: recovered
	0x82	0x20	Spoke-to-spoke time error: recovered

Sense Key	Sense Code	Sense Qualifier	Meaning	
0x02 Not Ready	0x04	0x00	Drive Not Ready	
	0x04	0x01	Drive Not Ready, in process of coming ready	
	0x04	0x02	Drive Not Ready, use START cmd	
	0x04	0x04	Drive Not Ready, formatting	
	0x04	0x09	Drive Not Ready, self test in progress	
	0x04	0x80	FileSystem Boot Area unreadable @ POR	
	0x04	0x82	File has not been written - can not read	
	0x04	0x83	Error reading diskware file	
	0x04	0x84	Diskware not loaded yet - issue start cmd	
	0x04	0x85	Code jumped to Dead Vector	
	0x04	0x86	Drive Not Ready, config page 8/HDA mismatched	
	0x04	0x88	Error reading Config Page file	
	0x04	0x91	Drive Not Ready, selfscan is running	
	0x04	0x92	NVR image not compatible	
	0x19	0x01	Drive Not Ready, Unknown defect list	
	0x19	0x03	Drive Not Ready, Bad grown defect list	
	0x31	0x00	Drive Not Ready, Format corrupted	
	0x40	0x80	Drive Not Ready, Unknown HW problem	
	0x40	0x81	Drive Not Ready, NVR is Corrupt	
	0x40	0x82	Drive Not Ready, CP section of EEPROM is Corrupt	
	0x40	0x83	Drive Not Ready, Failed HDA test	
	0x40	0x84	Drive Not Ready, Failed PCB test	
	0x40	0x85	Drive Not Ready, PM section of EEPROM is Corrupt	
	0x40	0x86	Drive Not Ready, Flash vector table is Corrupt	
	0x40	0x99	ROM and resident code are incompatible	
	0x03 Medium Error	0x02	0x00	No seek complete.
		0x03	0x01	Pre-Amp write unsafe
		0x03	0x84	DIAG: Write fault
0x0C		0x00	Write error	
0x0C		0x02	Write error - auto reallocation failed.	
0x0C		0x03	Write error - recommend reassignment	
0x11		0x00	Unrecovered read error	
0x11		0x04	Unrecovered read error - auto reallocate failed	
0x11		0x0C	Unrecovered read error - pending reallocation	
0x11		0x8B	Unrecovered read error - thermal asperity error	
0x11		0x81	System file has checksum error	
0x11		0x84	Error reading Diskware vector table file	
0x15		0x00	Positioning error detected by drive, check UPT	
0x15		0x02	Positioning error detected by read	
0x15		0x80	Positioning error detected by write	

Sense Key	Sense Code	Sense Qualifier	Meaning
0x03 Medium Error	0x15	0x81	Servo Write Not OK
	0x15	0x82	Servo Read Not OK
	0x15	0x83	Spoke address error
	0x16	0x00	Data synchronization mark error
	0x16	0x83	Spoke synchronization mark error
	0x19	0x00	MCS is invalid
	0x19	0x01	Unknown defect list
	0x19	0x03	Bad grown defect list
	0x19	0x80	Replacement block bad
	0x19	0x81	Found two bad RBNs
	0x27	0x00	Write protected
	0x31	0x00	Format corrupted
	0x31	0x01	Format command failed
	0x32	0x00	No available replacement block
	0x32	0x01	Failed replacement I/O (servo dead?).
	0x41	0x86	Required HW process did not complete
	0x44	0x80	Error writing FileSystem (neg cyl) file
	0x44	0x81	Error reading FileSystem (neg cyl) file.
	0x80	0x00	Forced error
	0x80	0x10	Block Slow PLL Error
	0x80	0x11	Split SLOW PLL Error
	0x80	0x12	Block state machine disabled error
	0x82	0x20	Spoke-to-spoke time error
	0x04 Hdwe Error	0x02	0x00
0x03		0x00	Write fault
0x06		0x80	PES calibration results are bad
0x06		0x81	BIAS calibration results are bad
0x06		0x82	ID GAMMA calibration values out of range
0x06		0x83	Recal record timeout
0x06		0x84	Seek fail during feed forward scale Cal.
0x06		0x85	Initial track grab failed.
0x06		0x86	Seek failure during Bias Calibration
0x06		0x87	Seek failure during KT/J Calibration
0x06		0x88	Seek failure during PES Calibration
0x06		0x89	Recal servo not ready
0x06		0x8A	Switch Point calibration results are bad
0x06		0x8B	Recal track number bad
0x06		0x8C	Seek failure during switch point cal
0x06		0x8D	Not all Head(s) in USER head map were found alive
0x06		0x8E	Valid SAMs not found!!!
0x06		0x8F	Sync To Disk Spokes failed

Sense Key	Sense Code	Sense Qualifier	Meaning
0x04 Hdwe Error	0x06	0x90	Invalid CP18 on disk
	0x06	0x91	Invalid CP21 on disk
	0x06	0x92	Invalid CP22 in NVR
	0x06	0x93	Head polarity cal failed.
	0x06	0x94	Svo: Invalid head polarity vector in CP23
	0x06	0x95	Seek 1 failure during servo writer erase band resync
	0x06	0x96	Seek failure during ARCs Calibration
	0x06	0x97	Seek failure during Pes Linearization
	0x06	0x98	Error during Pes Lin bode measurement
	0x06	0x99	Seek 2 failure during servo writer erase band resync
	0x06	0x9A	Seek failure during pes linearization table update
	0x06	0x9B	Motor speed error after turbo spin recovery
	0x06	0x9C	Electrical offset calibration returned an invalid value
	0x06	0x9D	Went servo dead during head polarity/coarse ATT cal.
	0x06	0x9E	Data consistency failure during PES Linearization
	0x06	0x9F	Couldn't find sync in the erase band & couldn't restart servo
	0x06	0xA0	Went servo dead after completing a seek
	0x06	0xA1	Successful sync not in the erase band
	0x06	0xA2	ID GAMMA projected value out of range
	0x06	0xA3	OD GAMMA calibration values out of range
	0x06	0xA4	OD GAMMA projected value out of range
	0x06	0xA5	Could not start estimator on printed media
	0x06	0xA6	Too many spoke errors on printed media-1st
	0x06	0xA7	Too many spoke errors on printed media-2nd
	0x06	0xA8	Too much variation in position at crash stop
	0x06	0xA9	Unable to determine a unique index
	0x06	0xAA	Too much runout at the crash stop even with BCVs
	0x06	0xAB	Crash stop position too large or too small
	0x06	0xAC	Too much position runout at the crash stop
	0x06	0xAD	Milli plant mag out of range
	0x06	0xAE	Milli loop mag out of range
	0x06	0xAF	Seek error for gamma frequency
	0x06	0xB0	Seek error for FF Gain
	0x06	0xB1	Cant find good track
	0x06	0xB2	Servo Busy too long
	0x09	0x00	Track following error
	0x09	0x91	Svo:TA Err-1st
	0x09	0x92	Svo:TA Err-2nd
	0x09	0x93	Svo:Bad Sync-1st
	0x09	0x94	Svo:Bad Sync-2nd
	0x09	0x95	Svo: Offtrack Prediction Err-1st

Sense Key	Sense Code	Sense Qualifier	Meaning
0x04 Hdwe Error	0x09	0x96	Svo: Offtrack Prediction Err-2nd.
	0x09	0x97	Svo:Head Spoke Number Miscompare-1st
	0x09	0x98	Svo:Head Spoke Number Miscompare-2nd
	0x09	0x99	Svo:Timeout waiting for burst/Spoke-1st
	0x09	0x9A	Svo:Timeout waiting for burst/Spoke-2nd
	0x09	0x9B	Svo: BCV err - 1st
	0x09	0x9C	Svo: BCV err - 2nd
	0x09	0x9D	Svo: Serial port timeout - 1st
	0x09	0x9E	Svo: Serial port timeout - 2nd
	0x09	0xA0	Svo:Possible Bump-1st
	0x09	0xA1	Svo:Bump Err-2nd
	0x09	0xA2	Svo: Timeout waiting for burst/Spoke - 1st
	0x09	0xA3	Svo: Timeout waiting for burst/Spoke - 2nd
	0x09	0xA4	Svo:Map'd Wdg-1st
	0x09	0xA5	Svo:Map'd Wdg-2nd
	0x09	0xA6	Svo:Id Err-1st.
	0x09	0xA7	Svo:Id Err-2nd
	0x09	0xA8	Svo:Unknown Wdg Err-1st
	0x09	0xA9	Svo:Unknown Wdg Err-2nd
	0x09	0xAA	Svo:Shock detected error-1st
	0x09	0xAB	Svo:Shock detected error-2nd
	0x09	0xAC	Svo: Seek time out
	0x09	0xAD	Svo: Bad measured timing values
	0x09	0xAE	Svo: Not spinning at proper speed-1st
	0x09	0xAF	Svo: Not spinning at proper speed-2nd
	0x09	0xB0	Svo:Confined sync timeout
	0x09	0xB1	Svo:R/W serial timeout
	0x09	0xB2	Svo:Context Acquire Timeout
	0x09	0xB3	Svo:Context Aquire Data Addr Valid Timeout
	0x09	0xB4	Svo:Context Aquire Bad Track Number
	0x09	0xB5	Svo: False SSMF - 1st
	0x09	0xB6	Svo: False SSMF - 2nd
	0x09	0xB7	Svo: Servo Dead Loading Channel
	0x09	0xB8	Svo: Servo Dead
	0x09	0xB9	Svo: Servo died during sync-to-index
	0x09	0xBA	Svo: Servo Dead Seek Command
	0x09	0xBB	Svo: Bogus seek command type
	0x09	0xBC	Svo: Servo Command Procesor Timeout
	0x15	0x00	Random positioning error
	0x15	0x02	Positioning error detected by read
0x32	0x00	A-list is full	

Sense Key	Sense Code	Sense Qualifier	Meaning
	0x3E	0x03	Logical Unit failed self test
	0x3E	0x04	Logical unit log update failure

Sense Key	Sense Code	Sense Qualifier	Meaning
0x04 Hdwe Error	0x40	0x80	Unknown HW problem
	0x40	0x82	Hardware Bug Check (check Info field)
	0x40	0x83	Failed HDA test.
	0x40	0x84	Failed PCB test.
	0x40	0x85	EDC check failed.
	0x40	0x86	Controller FIFO over/underflow error.
	0x40	0x88	DIAG: Drive failed to read servowriter serial number.
	0x40	0x89	DIAG: Failed XOR test.
	0x40	0x90	EDC Check Failed 0
	0x40	0x91	EDC Check Failed 1
	0x40	0x92	EDC Check Failed 2
	0x40	0x93	EDC Check Failed3
	0x40	0x8A	Servo restart failed after set RPM command
	0x40	0x8C	LBA mismatch error
	0x40	0x8D	ETM test failure, drive failed to spin
	0x40	0x8E	Could not achieve target RPM
	0x40	0xA0	Test did not complete
	0x40	0xA1	Test completed successfully, but failed.
	0x40	0xA2	Invalid parameter was received.
	0x40	0xA3	Invalid data in servo record. Prevented algorithm from completing.
	0x40	0xA4	Servo record did not complete
	0x40	0xA5	Read seek failed
	0x40	0xA6	Write seek failed
	0x40	0xA7	Raw seek failed
	0x40	0xA8	Servo Diagnostic Test running.
	0x40	0xA9	Invalid 97/98 super cmd seq. Must send 97 before 98
	0x40	0xAA	Super Command 97 already running
	0x40	0xAB	Complex magnitude is zero
	0x40	0xAC	Could not complete the low-level BCV write function
	0x40	0xAD	Got a divide by zero error in complex divide.
	0x40	0xAE	WGATE negated due to shock or bump during BCV write
	0x40	0xAF	Servo Death occurred during a BVC write operation
	0x40	0xB0	Seek command completion timeout occurred
	0x40	0xB1	Read an impossibility high number of good BCVS on this
	0x40	0xB2	No read passes resulted in enough test BCVs being read
	0x40	0xB3	The last read pass for test BCVs was good but shouldn't be
	0x40	0xB4	DestinationLoc expected vs actual mismatch. servo death recovery
	0x44	0x00	Logical assertion (firmware consistency check)
	0x44	0x84	Insufficient buffer memory for operation
	0x90	0x01	Serial Port Packet Receive Error

Sense Key	Sense Code	Sense Qualifier	Meaning
0x04 Hdwe Error	0xA0	0x00	New wrt cmd overlaps deferred write error
	0xA0	0x01	New wrt cmd overlaps deferred read error
	0xA0	0x02	New rd cmd overlaps deferred write error
	0xA0	0x03	New rd cmd overlaps deferred read error
	0xA0	0x81	Failed to load the FSBA
	0xA0	0x82	a power-up servo calibration failed
	0xA0	0x83	head check failed
	0xA0	0x84	Loading of disk config pages from filesystem failed
	0xA0	0x85	Diskware loaded, but failed validation
	0xA0	0x86	Loading of disware from filesystem failed
	0xA0	0x87	Validation of Vector Table failed
	0xA0	0x88	Failed to load diskware vector table from filesystem
	0xA0	0x89	Servo could not successfully complete pre-diskware calibrations
	0xA0	0x8A	Servo couldn't successfully complete post-diskware calibrations
	0xA0	0x8B	config page 19 corrupt
	0xA0	0x8C	No disk config pages loaded
	0xA0	0x8D	Loading of Orca table has failed
	0xA0	0x8E	Loading of log pages from file system has failed
	0xA0	0x8F	DBN write/read failed on head
	0xA0	0x90	Diagnostics failed on unit
	0xA0	0x91	Failed to spin up in time
	0xA0	0x92	Head load failed
	0xA0	0x93	FLT bit set in spinup
	0xA0	0x94	Selfscan could not load Script (0x30)

Sense Key	Sense Code	Sense Qualifier	Meaning
0x04 Hdwe Error	0xA0	0x95	Selfscan could not load results (0x31)
	0xA0	0x96	Selfscan could not load Defect (0x32)
	0xA0	0x97	Selfscan could not load Defect Manager files
	0xA0	0x98	Selfscan could not load SuperZoneData file
	0xA0	0x99	Selfscan timeout waiting for drive to come ready
	0xA0	0x9A	Load Selfscan overlays failed
	0xA0	0x9B	Selfscan overlays do not match firmware
	0xA0	0x9C	Failed to lock within TBD seconds0x04/0xA0/0xA0 Spin time-out INIT
	0xA0	0x9D	Retract timeout
	0xA0	0x9E	Drive Not Ready, CP section of EEPROM is Corrupt
	0xA0	0x9F	Drive Not Ready, PM section of EEPROM is Corrupt
	0xA0	0xA0	Spin time-out - INIT
	0xA0	0xA1	Spin time-out - STOP_PENDING
	0xA0	0xA2	Spin time-out - BRAKE
	0xA0	0xA3	Spin time-out - Stopped
	0xA0	0xA4	Spin time-out - PSENSE
	0xA0	0xA5	Spin time-out - DOUBLE ALIGN
	0xA0	0xA6	Spin time-out - OPEN_LOOP
	0xA0	0xA7	Spin time-out - SYNC
	0xA0	0xA8	Spin time-out - SYNC2
0xA0	0xA9	Spin time-out - CLOSED_LOOP	
0xA0	0xAA	Spin time-out - CLOSED_LOOP-2	
0xA0	0xAB	Spin time-out - CLOSED_LOOP-3	

Sense Key	Sense Code	Sense Qualifier	Meaning
0x04 Hdwe Error	0xA0	0xAC	Spin time-out - SETTLE
	0xA0	0xAD	Spin time-out - STEADY
	0xA0	0xAE	Spin time-out - FAULT
	0xA0	0xAF	Spin time-out - ETM
	0xA0	0xB0	Spin time-out - ROCKSPINDLE
	0xA0	0xB1	Spin time-out - ANTILOCK
	0xA0	0xB2	Spin time-out - SYNC4BRAKE
	0xA0	0xB3	Diablo shutdown due to high temp
	0xA0	0xB4	VCM wasn't updated in time
	0xA0	0xB5	Spindle clk is missing
	0xA0	0xB6	Sync-to-Index, servo went dead
	0xA0	0xB7	Sync-to-Index failed
	0xA0	0xB8	No spokes found
	0xA0	0xB9	Seeking too long
	0xA0	0xBA	Potential spindle out of synch
	0xA0	0xBB	fast head switch channel/preamp count mismatch
	0xA0	0xC0	Drive Not Ready, NVR G Image copies corrupt
	0xA0	0xC1	Drive Not Ready, PM section has invalid version id
	0xA0	0xC2	Context acquire timeout
	0xA0	0xC3	Track acquire timeout
	0xA0	0xC4	Head Pin Layer Reversal
	0xA0	0xC5	Heads cannot find servo sync (no PLR but heads with nosyncFound).
	0xA0	0xC6	Could not find sync in the erase band
	0xA0	0xC8	Could not arm Block State Machine, no matching pairs
	0xA0	0xC9	Could not arm Block State Machine, retry exhausted
	0xA0	0xCA	Could not arm Block State Machine
	0xA0	0xCB	OFM logical list synchronization timeout
	0xA0	0xCD	Bias Cal failed during power-up calibrations
	0xA0	0xCE	Gamma Cal failed during power-up calibrations
	0xA0	0xCF	ARCS Cal failed during power-up calibrations

Sense Key	Sense Code	Sense Qualifier	Meaning
0x05 Illegal Request	0x1A	0x00	Parameter list length error
	0x20	0x00	Invalid command.
	0x20	0xFE	ETM can't run with motor spinning
	0x20	0xFF	Selfscan code not loaded or bad
	0x21	0x00	Invalid LBA
	0x24	0x00	Invalid bits set in CDB
	0x25	0x00	Invalid LUN specified
	0x26	0x00	Invalid fields in parameters
	0x26	0x02	Bad fields in parameters
	0x26	0x04	Invalid release of persistent reservation
	0x26	0x80	Invalid Code image checksum failed
	0x26	0x81	Incorrect product code loaded
	0x26	0x82	Invalid Offset for image
	0x26	0x83	Invalid ASIC for this code load
	0x26	0x84	Invalid Code image - New PCBA cannot take old code
	0x2C	0x00	Invalid command sequence
	0x3D	0x00	Invalid bit set in identify message
	0x84	0x01	Aborted b/c of internal TO DER of a cmd other than current
	0x84	0x10	Aborted b/c of internal TO before going to BBR for cur cmd
0x84	0x11	Aborted b/c of internal TO BBR of a cmd other than current	

Sense Key	Sense Code	Sense Qualifier	Meaning
0x06 Unit Attn	0x0B	0x01	SMART overtemperature trip
	0x28	0x00	Not Ready to Ready transition
	0x29	0x01	Power-On-Reset occurred
	0x29	0x02	SCSI bus reset occurred
	0x29	0x03	Bus device reset message
	0x29	0x04	Recovered from bug check
	0x29	0x05	Bus mode changed to single ended
	0x29	0x06	Bus mode changed to LVDS
	0x29	0x86	Reset after too many ACK's
	0x2A	0x01	Mode parameters were changed
	0x2A	0x02	Log parameters were changed
	0x2A	0x03	Persistent Reservations were cleared
	0x2A	0x04	Persistent Reservations were released
	0x2A	0x05	Registration Preempted
	0x2f	0x00	Cmd Queue cleared by another initiator
	0x3f	0x01	Microcode has been changed
	0x3f	0x05	Device Identifier changed
	0x5C	0x01	Spindle synced
	0x5C	0x02	Spindle not synced
	0x06 Unit Attn	0x5D	0x00
0x5D		0x01	Smart triggered due to Write TA Factor
0x5D		0x02	Smart triggered due to Write Err Factor
0x07 Data Protect	0x5D	0x5B	Smart triggered due to spin retry
	0x5D	0x80	Smart test tripped
	0x5D	0xFF	Smart test tripped
0x09 Vendor Specific	0x27	0x00	Data write protected
	0x81	0x00	New command overlaps deferred write error
	0x81	0x01	New wrt cmd overlaps deferred write
	0x81	0x02	New rd cmd overlaps deferred write
	0x81	0x03	New rd cmd overlaps deferred read error

Sense Key	Sense Code	Sense Qualifier	Meaning
0x0B Aborted Cmd	0x00	0x00	Aborted.
	0x0E	0x01	Abort: IU too short (data frame with length = 0)
	0x1B	0x00	Abort: Synchronous transfers timeout
	0x25	0x00	Aborted with invalid LUN
	0x3F	0x0F	Aborted: Echo buffer overwritten
	0x43	0x00	Abort: Message reject error
	0x45	0x00	Abort: Reselection failed
	0x47	0x00	Abort: Bus parity error
	0x47	0x01	Abort: pCRC error detected
	0x47	0x02	Abort: Bus parity error during ST data phase
	0x47	0x03	Abort: Information Unit CRC error
	0x47	0x04	Abort: Async Protection error
	0x48	0x00	Abort: Initiator detected error
	0x49	0x00	Aborted with invalid Message
	0x4B	0x02	Abort: Too much write data received
	0x4B	0x03	Abort: ACK/NAK timeout
	0x4B	0x04	Abort: NAK received
0x4B	0x05	Abort: Data offset error	
0x4E	0x00	Overlapped commands attempted	
0x84	0x00	Aborted b/c of internal TO in DER for cur cmd	
0x0E	0x1D	0x00	Data miscompare error

Table 5-173 Sense-Key Specific Field Contents

Sense Key	Field Contents	Initiator Interpretation
ILLEGAL REQUEST	Field pointer bytes (Table 5-174)	Illegal field's offset (in the Command Descriptor Block or data-out buffer transfer parameters).
NOT READY	Progress Indication (Table 5-176)	Indicates progress of FORMAT UNIT Command.
RECOVERED ERROR	Retry count (Table 5-178)	Number of retries for I/O operation
MEDIUM ERROR	Retry count (Table 5-178)	Number of retries for I/O operation

The data format of the Field Pointer bytes associated with the ILLEGAL REQUEST sense key is shown in Table 5-174; the description of the data is in Table 5-175.

Table 5-174 ILLEGAL REQUEST Sense Key Field Pointer Bytes/Data Format

Bit Byte	7	6	5	4	3	2	1	0
15	SKSV = 1	C/D	Reserved		BPV	Bit Pointer		

16 - 17	Field Pointer
---------	---------------

Table 5-175 *ILLEGAL REQUEST Sense Key Field Pointer Bytes\Field Descriptions*

Field	Description
SKSV	Sense-Key Specific Valid
C/D	Command/Data. When the value in this field is 1, the Field Pointer identifies the first byte in error in the Command Descriptor Block. When 0, the Field Pointer identifies the first byte in error in the data parameters sent by the initiator in the data-out buffer transfer.
BPV	Bit Pointer Valid. The drive does not support Bit-level resolution; the bit value must be 0.
Bit Pointer	The drive does not support Bit-level resolution; the bit value must be 0.
Field Pointer	Indicates the first byte in error in either the Command Descriptor Block or in the data parameters sent by the initiator.

The data format of the Progress Indication bytes associated with the NOT READY sense key is shown in Table 5-176; the description of the data is contained in Table 5-177.

Table 5-176 *NOT READY Sense Key - Progress Indication Bytes\Data Format*

Bit Byte	7	6	5	4	3	2	1	0
15	SKSV = 1	Reserved						
16 - 17	Progress Indication							

Table 5-177 *NOT READY Sense Key - Progress Indication Bytes\Field Descriptions*

Field	Description
SKSV	Sense-Key Specific Valid
Progress Indication	Indication of percent complete for the FORMAT UNIT command. The returned value is the numerator of a fraction that has 65536 (10000h) as the denominator (Progress Indication/10000h = percent complete).

The data format of the Retry Count bytes associated with the MEDIUM ERROR or RECOVERED ERROR sense key is shown in Table 5-178; the description of the data is contained in Table 5-179.

Table 5-178 *MEDIUM ERROR or RECOVERED ERROR Sense Key - Retry Count-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
15	SKSV = 1	Reserved						
16 - 17	Retry Count							

Table 5-179 *MEDIUM ERROR or RECOVERED ERROR Sense Key - Retry Count*

-Field Descriptions

Field	Description
SKSV	Sense-Key Specific Valid
Retry Count	The number of times an I/O operation was retried

5.33 RESERVE (6) Command (16h)

The RESERVE and the RELEASE commands are used for contention resolution in multiple-initiator systems. The RESERVE command is used to reserve a logical unit. The data format and the field descriptions are described in Table 5-180 and Table 5-181 respectively.

Table 5-180 RESERVE (6) Command Descriptor Block-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (16h)							
1	Reserved			3rdPty	Third Party Device ID			Extent
2	Reservation Identification							
3 - 4	Extent List Length							
5	Control							

Table 5-181 RESERVE (6) Command-Field Descriptions

Field	Description
3rd Pty	When set to 1, it indicates that the Third Party Device ID field is valid.
Third Party Device ID	<p>Required and used only when the 3rdPty bit is set, in which case this field specifies the SCSI ID (0 through 7 only) of the initiator to be granted the reservation of the logical unit. The drive preserves the reservation until one of the following occurs:</p> <ul style="list-style-type: none"> • It is superseded by another valid RESERVE command from the initiator. • It is released by the same initiator • It is released by a TARGET RESET message from any initiator. • It is released by a hard reset condition. <p>The drive ignores any attempt to release the reservation made by any other initiator. For example, if ID7 sends ID2 a Third Party reservation on behalf of ID6 (the target at ID2 gets reserved for the initiator ID6), then only ID7 can release the target at ID2 (using a Third Party release). ID6 cannot release the reservation even though the reservation was made on its behalf.</p>
Extent	The drive supports reservations only on entire logical units. The value must be 0.
Reservation Identification	Any value in this field is ignored by the drive.
Extent List Length	Any value in this field is ignored by the drive.

Note: The drive returns RESERVATION CONFLICT status if the unit is already reserved for another initiator. Once the reservation is honored, any command from another initiator (except INQUIRY or REQUEST SENSE) is rejected with RESERVATION CONFLICT status.

5.34 RESERVE (10) Command (56h)

The RESERVE and the RELEASE commands are used for contention resolution in multiple-initiator systems. The RESERVE command is used to reserve a logical unit. If RESERVE (10) is used, then RELEASE (10) is also used. The data format and the field descriptions are described in Table 5-182 and Table 5-183 respectively. The data format for the Event Descriptors is described in Table 5-184.

Table 5-182 RESERVE (10) Command Descriptor Block-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (56h)							
1	Reserved			3rdPty	Reserved		LongID	Extent
2	Reservation Identification							
3	Third Party Device ID							
4 - 6	Reserved							
7 - 8	(MSB)	Parameter List Length						(LSB)
9	Control							

Table 5-183 RESERVE (10) Command-Field Descriptions

Field	Description
3rd Pty	When set to 1, it indicates that the Third Party Device ID field is valid.
LongID	When 0, indicates that the third party device associated with the reservation release has a number smaller than 255 and the ID value can be sent within the CDB. If set = 1, indicates that the third party device ID is greater than 255, the ID value within the CDB is ignored, and the parameter list length is at least eight.
Third Party Device ID	<p>Required and used only when the 3rdPty bit is set, in which case this field specifies the SCSI ID of the initiator to be granted the reservation of the logical unit. The drive preserves the reservation until one of the following occurs:</p> <ul style="list-style-type: none"> • It is superseded by another valid RESERVE command from the initiator. • It is released by the same initiator. • It is released by a TARGET RESET message from any initiator. • It is released by a hard reset condition. <p>The drive ignores any attempt to release the reservation made by any other initiator. For example, if ID7 sends ID2 a Third Party reservation on behalf of ID6 (the target at ID2 gets reserved for the initiator ID6), then only ID7 can release the target at ID2 (using a Third Party release). ID6 cannot release the reservation even though the reservation was made on its behalf.</p>
Extent	<p>If 0, requests that the entire logical unit be reserved for use only by the initiator until the request is supplanted by another valid RESERVE command from the same initiator, or until released via a RELEASE (10) command from the reserving initiator, a hard reset, or a power on cycle.</p> <p>If set = 1, the extent reservation option is implemented. This option allows an application client in a multitasking environment to have multiple reservations.</p> <p>The size of the extent list is defined by the contents of the Parameter List Length field. The extent list consists of zero or more descriptors. Each descriptor defines an extent beginning at the specified logical block address for the specified number of blocks. If the number of blocks is 0, the extent begins at the specified logical block address and continues through the last logical block address on the logical unit. The data format of extent descriptors is shown in Table 5-184.</p>

Table 5-184 *Extent Descriptors-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0	Reserved					RelAdr	Reservation Type	
1 - 3	(MSB) Number of Blocks							(LSB)
4 - 7	(MSB) Logical Block Address							(LSB)

The Reservation Type field defines the type of reservation in effect for the extent being reserved. Available types of reservations are shown in Table 5-185.

Table 5-185 *Reservation Types*

Reservation Type	Description
00b	READ shared. No write operations are permitted by any initiator to the reserved extent. Any initiator may read from the extent.
01b	WRITE exclusive. No other initiator may write to the indicated extent. Any initiator may read from the extent.
10b	READ exclusive. While this reservation is active, no other initiator may be permitted to read to the indicated extent.
11b	Exclusive access. While this reservation is active, no other initiator is permitted any access to the reserved extent.

RelAdr is always 0: The Maxtor Atlas 10K V hard disk drive does not support relative addressing.

If both the LongID and the Extent bits = 1, then the parameter list takes the format shown in Table 5-186.

Table 5-186 *Parameter List When LongID and Extent Bits = 1-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0 - 7	(MSB) Third Party Device ID							(LSB)
8 - n	(MSB) Extent Descriptors							(See Figure 5-91 Above)(LSB)

If the LongID bit = 1 and the Extent bit = 0, then the parameter list length is eight and the parameter list has the format shown in Table 5-187.

Table 5-187 *RESERVE (10) ID Only Parameter List-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0 - 7	(MSB) Third Party Device ID							(LSB)

5.35 REZERO UNIT Command (01h)

The REZERO UNIT command sends the heads to cylinder 0 (Logical Block Address 0). Once the heads are on the track, they are not guaranteed to stay at cylinder 0; periodic calibrations may send them to a different cylinder. If periodic calibrations have been disabled by using the MAN bit (byte 3, bit 6) of the Vendor Unique Function Control Page in the MODE SELECT command (15h), the REZERO UNIT command will cause the drive to perform the next scheduled calibration. The data format is shown in Table 5-188.

Table 5-188 REZERO UNIT Command Descriptor Block-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (01h)							
1 - 4	Reserved							
5	Control							

5.36 SEEK (6) Command (0Bh)

The SEEK (6) command uses a 6-byte Command Descriptor Block to request the drive to send the drive heads to the requested Logical Block Address. The drive responds by placing the heads at the correct Logical Block Address. When the heads are at the specified Logical Block Address, they are not guaranteed to stay in that position since periodic calibrations may send them to a different cylinder. The data format is shown in Table 5-189.

Table 5-189 *SEEK (6) Command Descriptor Block-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (0Bh)							
1	Reserved			Logical Block Address				
2 - 3	Logical Block Address							
4	Reserved							
5	Control							

5.37 SEEK (10) Command (2Bh)

The SEEK (10) command uses a 10-byte Command Descriptor Block to request the drive to send the drive heads to the requested Logical Block Address. The drive responds by placing the heads at the correct Logical Block Address. When the heads are at the specified Logical Block Address, they are not guaranteed to stay in that position since periodic calibrations may send them to a different cylinder. The data format is shown in Table 5-190.

Table 5-190 *SEEK (10) Command Descriptor Block-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (2Bh)							
1	Reserved							
2 - 5	Logical Block Address							
6 - 8	Reserved							
9	Control							

5.38 SEND DIAGNOSTIC Command (1Dh)

The SEND DIAGNOSTIC command sends a diagnostic page to the drive. Refer to the RECEIVE DIAGNOSTICS RESULTS command for a description of the diagnostic pages and the use of the diagnostic commands. The Supported Diagnostic Page List (Table 5-193) and Translate Page List (Table 5-194) can be sent with this command. Table 5-191 shows the SEND DIAGNOSTIC CDB data format and Table 5-192 shows the field descriptions.

Table 5-191 SEND DIAGNOSTIC Command Descriptor Block-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (1Dh)							
1	Function			PF	Rsvd	Self Test	DevOnl	UnitOfl
2	Reserved							
3 - 4	Parameter List Length							
5	Control							

Table 5-192 SEND DIAGNOSTIC Command-Field Descriptions

Field	Description
Function	<p>When the SelfTest bit is set to one, the Function code field is 000b. When the SelfTest bit is cleared to zero, the content of the Function code field is specified as follows:</p> <p>000b: Value to be used when the SelfTest bit is set to one or if the SEND DIAGNOSTIC command is not invoking one of the other self-test function codes (for example, Translate Address page)</p> <p>001b: Background short self-test. The device server will start its short self-test routine in the background mode.</p> <p>010b: Background extended self test. The device server will start its extended self-test routine in the background mode.</p> <p>011b: Reserved</p> <p>100b: Abort background self-test. Abort the current self-test running in background mode. This value is only valid if a previous SEND DIAGNOSTIC command specified a Background self-test function and that function has not completed. If either of these conditions is not true, then the device server will respond by returning a CHECK CONDITION with a sense key of NOT READY and an additional sense code of LOGICAL UNIT NOT READY, SELF-TEST IN PROGRESS.</p> <p>101b: Foreground short self-test. The device server will start its short self-test routine in the foreground mode</p> <p>110b: Foreground extended self-test. The device server will start its extended self-test routine in the foreground mode</p> <p>111b: Reserved</p>
PF	<p>Page Format. When the PF bit is 1, the parameters that follow conform to the Supported Diagnostic Page List (Table 5-193) or the Translate Address Page (Table 5-194).</p> <p>SelfTest is ignored when PF = 1.</p> <p>The PF bit must be 1 if Parameter List Length is not 0.</p>

Field	Description
SelfTest	If the PF bit is 0 and the SelfTest bit is 1, the drive runs its self-test routine. GOOD status is returned to the initiator if the self-test passes. CHECK CONDITION status is returned along with any available sense data if the self-test fails. This field is ignored if the PF bit is 1.
DevOnl	Any value in this field is ignored by the drive.
UnitOfl	Any value in this field is ignored by the drive.
Parameter List Length	When the Parameter List Length field is non-zero, the value specifies the length (in bytes) of the parameter list that is to be transferred from the initiator to the drive during the command's data-out buffer transfer. This field is only valid when the PF bit = 1 and must be 0 when the PF bit is 0.

5.38.1 Supported Diagnostic Page List

Table 5-193 shows the Supported Diagnostic Page List that can be supplied with a SEND DIAGNOSTIC command to request that the Supported Diagnostic Page List be returned after the next RECEIVE DIAGNOSTIC RESULTS command.

Table 5-193 Supported Diagnostic Page List-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0	Page Code (00h)							
1	Reserved							
2 - 3	Page Length (00h)							

5.38.2 Translate Address Page

Table 5-194 shows the Translate Address Page that can be supplied with a SEND DIAGNOSTIC command to request that the Translate Address Page be returned after the next RECEIVE DIAGNOSTIC RESULTS command.

Table 5-194 Translate Address Page-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0	Page Code (40h)							
1	Reserved							
2 - 3	Page Length (0Ah)							
4	Reserved					Supplied Format 000b (LBA)		
5	Reserved					Translate Format 101b (Physical Sector)		
6 - 9	Address to Translate (LBA Format)							
10 - 13	Reserved							

5.39 SET DEVICE IDENTIFIER Command (A4h)

The SET DEVICE IDENTIFIER command requests that the device identifier information in the logical unit be set to the value send via the SET DEVICE IDENTIFIER command's parameter list. The data format and field descriptions are shown in Table 5-195 and Table 5-196 respectively.

Upon successful completion of a SET DEVICE IDENTIFIER command, a Unit Attention is generated for all initiators except the one that issued the command.

Table 5-195 SET DEVICE IDENTIFIER Command Descriptor Block-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (A4h)							
1	Reserved			Service Action (06h)				
2 – 5	Reserved							
6 – 9	(MSB) Parameter List Length (LSB)							
10	Reserved							
11	Control							

Table 5-196 SET DEVICE IDENTIFIER Command Descriptor Block-Field Descriptions

Field	Description
Service Action	Must = 06h. Any other value forces Check Condition, Illegal Request.
Parameter List Length	This field specifies the length, in bytes, of the Identifier to be transferred from the application client to the device server. The maximum value for this field is 64 bytes. A parameter list length of 0 indicates that no data will be transferred, and that subsequent REPORT DEVICE IDENTIFIER commands will return an Identifier length of 0. If the parameter list length exceeds 64 bytes, then the drive returns a Check Condition status with the sense key set to Illegal Request, and an additional sense code of Invalid Field in CDB.

The SET DEVICE IDENTIFIER command's parameter list contains Identifier to be set by the logical unit to which the command is sent. The parameter list is shown below in Table 5-197 and the field descriptions are shown in Table 5-198.

Table 5-197 SET DEVICE IDENTIFIER Parameter List-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0 – n	(MSB) Identifier (LSB)							

Table 5-198 SET DEVICE IDENTIFIER Parameter List-Field Descriptions

Field	Description
Identifier	The value reported in this field is the data to be returned by all subsequent REPORT DEVICE IDENTIFIER commands, until a new SET DEVICE IDENTIFIER command is issued that changes the data.

5.40 START STOP UNIT Command (1Bh)

The START STOP UNIT command enables or disables the logical unit for media access operations. This command gives the initiator remote control over drive spin-up and spin-down. The data format and field descriptions are shown in Table 5-199 and Table 5-200 respectively.

A jumper (or backplane connection) is used to select the spin-up mode. When configured for remote control of spin-up (Option 3), the drive requires a START STOP UNIT Command to spin up. See Chapter 3 for jumper-programming instructions for your specific hard disk drive.

Table 5-199 *START STOP UNIT Command Descriptor Block-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (1Bh)							
1	Reserved							Immed
2 - 3	Reserved							
4	Power Conditions				Reserved		LoEj	Start
5	Control							

Table 5-200 *START STOP UNIT Command-Field Descriptions*

Field	Description
Immed	Immediate. When 1, the drive returns GOOD status immediately after validating the Command Descriptor Block, but before the drive spins up. When 0, the drive does not return GOOD status until after the drive has spun up successfully and is ready for media access commands.
LoEj	Load Eject. This value must be zero, because the disk drive contains non-removable medium.
Start	When Start is 1, the unit spins up (if it is not currently spinning) and is left in a state where subsequent media access commands can succeed. When the value is 0, the drive spins down. Subsequent media access commands are rejected with CHECK CONDITION status and the proper sense data.
Power Conditions	The Power Conditions field is not supported.

5.41 SYNCHRONIZE CACHE Command (35h)

The SYNCHRONIZE CACHE command ensures that all logical blocks, within the specified range in the drive's cache memory, have their most recent data value recorded on the physical medium. If more recent data for a logical block, within the specified range, exists in the cache memory than on the physical medium, then the logical blocks from the cache memory is written to the physical medium. Logical blocks are not necessarily removed from the cache as a result of the Synchronize Cache operation. The data format and field descriptions are shown in Table 5-201 and Table 5-202 respectively.

Table 5-201 SYNCHRONIZE CACHE Command Descriptor Block-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (35h)							
1	Reserved						Immed = 0	RelAdr
2 - 5	Logical Block Address							
6	Reserved							
7 - 8	Number of Blocks							
9	Control							

Table 5-202 SYNCHRONIZE CACHE Command-Field Descriptions

Field	Description
Immed	Immediate. Must be 0. The drive returns status after the synchronize cache operation has been completed.
RelAdr	Relative Address. Must be 0. Relative addressing is not supported.
Logical Block Address	The address of the first logical sector of the range to be operated upon by this command.
Number of Blocks	Specifies the total number of contiguous logical blocks within the range. When set to 0, the drive removes all sectors from the given Logical Block Address on. When Number of Blocks and Logical Block Address are both 0, the drive removes all unwritten cache sectors to the medium. It is not considered an error when no sectors within the specified range are in the cache memory.

5.42 TEST UNIT READY Command (00h)

The TEST UNIT READY command returns GOOD status if the drive is ready to receive medium access commands (that is, READ, WRITE, VERIFY, etc.). If the drive is not ready to receive medium access commands, it returns CHECK CONDITION status with the appropriate sense keys and qualifiers. The data format is described in Table 5-203.

Table 5-203 TEST UNIT READY Command Descriptor Block-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (00h)							
1 - 4	Reserved							
5	Control							

5.43 VERIFY Command (2Fh)

The VERIFY command verifies the data on the medium for the Logical Block Address range specified. The data format and field descriptions are shown in Table 5-204 and Table 5-205 respectively.

Table 5-204 VERIFY Command Descriptor Block-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (2Fh)							
1	Reserved			DPO	Reserved		BytChk	RelAdr
2 - 5	Logical Block Address							
6	Reserved							
7 - 8	Verification Length							
9	Control							

Table 5-205 VERIFY Command-Field Descriptions

Field	Description
DPO	Disable Page Out. When the value of the field = 1, the device server will assign the logical blocks accessed by this command the lowest priority for being fetched into or retained by the cache. When DPO = 0, the priority will be determined by fields in the Caching page.
BytChk	Byte Check. When the value is 1, the drive performs a byte-by-byte comparison of the data on the medium with the data sent by the initiator in the data-out buffer transfer. When 0, the drive performs a medium verification with no data comparison. In this case, no data is sent by the initiator.
RelAdr	Relative Address. This feature is not supported; <i>the field must be 0.</i>
Logical Block Address	The address of the first logical sector to be verified.
Verification Length	The number of contiguous sectors that are to be verified. The drive returns a CHECK CONDITION status with the correct sense data if the specified transfer extends beyond out of the logical sector space.

5.44 WRITE (6) Command (0Ah)

The WRITE (6) command requests that the drive write initiator-supplied data to the medium.

The Command Descriptor Block for the WRITE (6) command is shown in Table 5-206; Table 5-207 describes the data fields.

Table 5-206 *WRITE (6) Command Descriptor Block-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (0Ah)							
1	Reserved			Logical Block Address				
2 - 3	Logical Block Address							
4	Transfer Length							
5	Control							

Table 5-207 *WRITE (6) Command-Field Descriptions*

Field	Description
Logical Block Address	The address of the first logical block where the Write operation will occur.
Transfer Length	The number of contiguous logical blocks of data to transfer. A transfer length of 0 indicates 256 logical blocks are to be transferred. Any other value indicates the actual number of logical blocks that will be transferred.

5.45 WRITE (10) Command (2Ah)

The WRITE (10) command requests that the drive write initiator-supplied data to the medium. The data format and field descriptions are shown in Table 5-208 and Table 5-209 respectively.

Table 5-208 *WRITE (10) Command Descriptor Block-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (2Ah)							
1	Reserved			DPO	FUA	Reserved		RelAdr
2 - 5	Logical Block Address							
6	Reserved							
7 - 8	Transfer Length							
9	Control							

Table 5-209 *WRITE (10) Command-Field Descriptions*

Field	Description
DPO	<p>Disable Page Out. When the value of the field is 1, the drive tells the controller to bypass cache as much as possible. The buffers used for the WRITE data are given a low priority. This bit may be useful if the host CPU is performing a background copy operation and does not want to wipe out the drive's cache.</p> <p>When DPO is 0, the drive tells the controller to cache this WRITE data in a normal manner.</p> <p>NOTE: The drive normally assigns a low priority to write caching. Therefore, the DPO bit in a WRITE command is ignored.</p>
FUA	<p>Force Unit Access. When 1, the data must be written on the medium before GOOD status is returned. When 0, GOOD status may be returned after the data is fetched from the initiator and before it is written to the medium. If the WCE (Write Cache Enable) bit on the MODE SELECT command Caching Page is clear, the FUA bit is ignored and behaves as if it were always set (that is, if WCE = 0, the controller forces FUA = 1).</p>
Logical Block Address	The address of the first logical block where the Write operation will occur.
RelAdr	Relative Address. This drive does not support this feature; the field must be 0.
Transfer Length	The number of contiguous logical blocks of data to write. The drive returns CHECK CONDITION status with the correct sense data if the specified transfer would cause the write operation to extend beyond the Logical Block Address space.

5.46 WRITE AND VERIFY Command (2Eh)

The WRITE AND VERIFY command writes host-supplied data to the unit and verifies it was written correctly. The data format and field descriptions are shown in Table 5-210 and Table 5-211 respectively.

Table 5-210 *WRITE AND VERIFY Command Descriptor Block-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (2Eh)							
1	Reserved			DPO	Reserved		BytChk	RelAdr
2 - 5	Logical Block Address							
6	Reserved							
7 - 8	Transfer Length							
9	Control							

Table 5-211 *WRITE AND VERIFY Command-Field Descriptions*

Field	Description
DPO	<p>Disable Page Out. When the value of the field is 1, the drive tells the controller to bypass cache as much as possible. The buffers used for the WRITE data are given a low priority. This bit may be useful if the host CPU is performing a background copy operation and does not want to wipe out the drive's cache.</p> <p>When DPO is 0, the drive tells the controller to cache this WRITE data in a normal manner.</p> <p>NOTE: The drive normally assigns a low priority to write caching. Therefore, the DPO bit in a WRITE command is essentially ignored.</p>
BytChk	<p>Byte Check. When 1, the drive performs a byte-by-byte comparison of the data written on the media with the data transferred from the initiator.</p> <p>When 0, the drive performs a medium verification with no data comparison of the sectors written by the command.</p>
RelAdr	Relative Address. This feature is not supported and the field must be 0.
Logical Block Address	The address of the first logical block where the WRITE operation will occur.
Transfer Length	The number of contiguous logical blocks of data to write and verify. The drive returns CHECK CONDITION status with the correct sense data if the specified transfer would cause the write operation to extend beyond the Logical Block Address space.

5.47 WRITE BUFFER Command (3Bh)

The WRITE BUFFER and the READ BUFFER commands allow the initiator to gain access to the data storage areas within the drive controller. These commands are used to check the integrity of the buffer-initiator data path independent of the storage medium, typically for fault isolation. The commands can be used to verify the caching policies of the controller or other kinds of debug work. In addition, the WRITE BUFFER command supports controller microcode downloads. The data format and field descriptions are shown in Table 5-212 and Table 5-213 respectively.

Both the READ BUFFER and WRITE BUFFER commands have a *Buffer ID* and a *Buffer Offset* data field defined in their Command Descriptor Block. In combined Header and Data Mode, if both are set to 0, reference is made to a single, 512-byte buffer set aside in the controller's data memory, independent of cache. This allows an initiator to verify the data path from the controller's data memory to the initiator without affecting the drive's internal data caches.

Table 5-212 WRITE BUFFER Command Descriptor Block-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (3Bh)							
1	Reserved				Mode			
2	Buffer ID							
3 - 5	Buffer Offset							
6 - 8	Parameter List Length							
9	Control							

Table 5-213 *WRITE BUFFER Command-Field Descriptions*

Field	Description
Mode	Five modes are supported by the drive; Combined Header and Data, Data Only, Download Microcode and Save, and Write Data from Echo Buffer. The permissible values for Buffer ID, Buffer Offset are as follows:
0000b	<p>Combined Header and Data – In this mode, the data to be transferred is preceded by a 4-byte header that contains reserved bytes that must be set to 0. The resulting descriptor and data are fetched from the initiator during the command's data-out buffer transfer and, after the descriptor has been validated, only the data is written to the controller's memory.</p> <p><i>Buffer ID</i> must be 0</p> <p><i>Buffer Offset</i> must be 0</p> <p><i>Parameter List Length</i> must be set to between 0 and 516 (4 bytes of header and up to 512 bytes of data).</p>
0010b	<p>Data Only– In this mode, the data-out buffer transfer contains buffer data. Data is written to the drive's buffer starting at the location specified by the Buffer ID and Buffer Offset.</p> <p><i>Buffer ID.</i> When non-zero, the value specifies a specific track line in the controller's cache. If Buffer ID of zero is specified, the buffer referenced is the same as if Buffer ID of one were specified.</p> <p><i>Buffer Offset.</i> When non-zero, the value can range from 0 to one less than the number of sectors/track multiplied by 512. This value must be either a zero or a multiple of 512.</p> <p><i>Parameter List Length.</i> This value must be a multiple of 512. The maximum value for this field is one more than the number of sectors/track times 512. The minimum value is 0.</p>
0101b	<p><i>Download Microcode and Save.</i> In this mode, a microcode image is transferred to the drive, and if valid, is saved in the controller's non-volatile memory. This image must be transferred with all buffers of the same size (the last buffer of the transfer cannot be of a different size) and the buffer size must be a multiple of 512. Most images are multiples of 512 and 1K. In that case, it is safe to have 512 byte transfers or 1K transfers. For example, if the image size is 369K, then 738 512-byte buffers can be transferred or 369 1K byte buffers can be transferred. It is also possible to transfer the image in one large buffer.</p> <p>The controller does a hard reset on command completion. In Download Microcode and Save mode, the other Command Descriptor Blocks are supported as follows:</p> <p><i>Buffer ID</i> – Undefined and ignored.</p> <p><i>Buffer Offset</i> – The offset of the buffer being transferred. For example: for 512 byte transfers use 0, 512, 1024, 1536,.... For 1K byte transfers use 0, 1024, 2048, 3072...</p> <p><i>Parameter List Length</i> – Must be the size of the buffer being transferred.</p>
7 (0111b)	<i>Download Microcode and Save with Offset.</i> This mode acts the same as mode 0101b but is restricted to sending the bytes in sequential order.
1010b	<p>Write Data from Echo Buffer</p> <p>In this mode, the host transfers data and stores it in an echo buffer. The data will be preserved in the echo buffer unless there is an intervening command.</p> <p><i>Buffer ID</i> Ignored.</p> <p><i>Buffer Offset</i> Ignored.</p>

5.48 WRITE LONG Command (3Fh)

The WRITE LONG and READ LONG commands allow the initiator to issue raw reads and writes of physical sectors addressed as logical blocks with an assumed sector length of 512 bytes. These commands are used to verify the correction capability of the drive's ECC and to create defects of a known kind and location to verify the drive's recovery and bad sector handling policies. The data format and field descriptions are shown in Table 5-214 and Table 5-215 respectively.

Table 5-214 WRITE LONG Command Descriptor Block-Data Format

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (3Fh)							
1	Reserved							RelAdr
2 - 5	Logical Block Address							
6	Reserved							
7 - 8	Byte Transfer Length (578)							
9	Control							

Table 5-215 WRITE LONG Command-Field Descriptions

Field	Description
RelAdr	Relative Address. This drive does not support this feature; the field must be 0.
Logical Block Address	The address of the first logical sector where the Write operation will occur.
Byte Transfer Length	<p>Must be 578. This value reflects the number of bytes spanned by the ECC, including the ECC itself, rounded up to the next word (this is easiest understood in bits: 4096 data bits + 1 force error bit + 15 Logical Block Address bits + 16 cross check bits + 2 pad bits + 480 ECC bits + 6 pad bits that are not written to media but are only needed for transfer = 4496 bits or 578 bytes). Any other value in this field results in a CHECK CONDITION status and the appropriate sense data.</p> <p>The initiator should first issue a READ LONG, then modify the data portion only, and then issue a WRITE LONG. It is also recommended that the drive be reformatted after these commands are used.</p>

5.49 WRITE SAME COMMAND

The WRITE SAME command requests that the drive write the single sector of data transferred by the initiator to the medium multiple times. The data format and field descriptions are shown in Table 5-11 and 5-12 Table respectively.

Table 5-216 *WRITE SAME Command Descriptor Block-Data Format*

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (41h)							
1	Reserved				PBdata	LBdata	RelAdr	
2-5	Logical Block Address							
6	Reserved							
7-8	Number of Blocks							
9	Control							

Table 5-217 *WRITE SAME Field Description*

Field	Description
PBdata	Physical Block Data - Not supported. This feild must be zero.
LBdata	Logical Block Data - A value of 1 requests that the drive replace the first four bytes of the data to be written to the current logical sector with the logical sector address of the sector currently being written.
RelAdr	Relative Address. - Not supported. This feild must be zero
Logical Block Address	The address of the first logical sector where the Write operation will occur.
Number of Blocks	This feild contains the number of contiguous logical blocks that are to be written. A value of 0 requests that all the remaining logical blocks on the medium are to be written.

5.50 WRITE SKIP MASK Command (EAh)

The WRTIE SKIP MASK command allows the initiator to condition the function of the WRITE commands to allow only selected data to be transferred over the SCSI interface. A WRITE SKIP MASK command precedes a WRITE (10) command; the WRITE (10) command must be linked to the WRITE SKIP MASK command or the skip mask is lost and a CHECK CONDITION status is returned.

Table 5-218 *WRITE SKIP MASK Command (EAh)*

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (EAh)							
1	Reserved							
2 - 5	(MSB)	Logical Block Address						(LSB)
6	Skip Mask Length							
7 - 8	(MSB)	Transfer Length						(LSB)
9	VU	Reserved				Flag	Link	

Table 5-219 *WRITE SKIP MASK Command - Field Description*

Field	Description
Logical Block Address	This field specifies the first logical block of the corresponding WRITE (10) COMMAND. This allows a SEEK to begin immediately. The LBA of the following WRITE (10) command must be the same as the LBA of the Write SKIP MASK command,
Skip Mask Length	This field indicated the length in bytes of the skip mask. The mask is up to 64 bytes, giving a data address capability of 256K. The mask consists of a sequence of bits where a "1" bit is for the blocks of data to be transferred and a "0" for the blocks of data that are to be skipped. The first "1" bit of the mask corresponds to the first LBA to be transferred. Any unused bits at the end of the last byte in the mask must be set to 0. The mask is transferred in the data-out buffer transfer of the skip mask command.
Transfer Length	This field must be equal to the transfer length in the linked WRITE (10) command. The transfer length field must be equal to the number "1" bits in the skip mask. If a WRITE (10) is used with 0 as the transfer length, a transfer length of 256 data blocks is implied.

Chapter 6

FEATURE DESCRIPTIONS

This chapter provides descriptions of many of the features incorporated in the Maxtor Atlas 10K V hard disk drive family.

6.1 KEY FEATURES

Maxtor Atlas 10K V disk drives provide unequaled quality and performance. They are designed to improve the performance of high-end workstations, file servers, and storage systems. Among their features are the following:

- 73.5, 147.1 GB, or 300 GB formatted capacity (512 bytes to 524 bytes/block selectable in increments of 2 bytes)
- GMR heads and banded recording for higher overall bit-densities
- Self-Monitoring, Analysis, and Reporting Technology (S.M.A.R.T.) Revision 2 support
- Optimized Reordering Command Algorithm and tagged command queuing for greater throughput
- 8 MB segmented cache buffer ECC protected
- 10,000 rpm rotational speed yielding 3 ms average latency
- 4.0 ms / 4.5 ms average read/write seek time for the 73.5 GB disk drives, 4.2 ms / 4.7 ms average read/write seek time for the 147.1 GB disk drives; and 4.4 ms / 4.9 ms average read/write seek time for the 300 GB disk drives.
- Embedded servo system for exceptional head positioning accuracy and long life
- “Zero latency” reads and writes reduce data access time
- Full SCSI-3 compliance¹ (compatible with SCSI-2 and SCSI-1)
- Variable BPI
- Just-in Time Seeks
- Read/Write Ahead

- MaxAdpat Feature Set
 - Rotational Vibration Compensation (RVC)
 - Adaptive Bias Estimation (ABE)
 - Virtual Cache Line (VCL)
 - Adaptive Active Filtration (AAF)
- Advanced Ultra160 SCSI-LVD, Ultra320 SCSI-LVD, Ultra2-LVD, and Ultra-SE with 68-pin wide connector or 80-pin SCA-2 connector
- Superior data integrity:
 - 352-bit Reed-Solomon quad-burst ECC provides fast on-the-fly correction for small errors and robust software-based correction for larger errors
 - EDC internal protection from SCSI bus to media

Note: ¹The drives are in compliance with the SCSI-3 draft proposed ANSI standards as they exist at the time of product release; SCSI-3 features may not be enabled on drives when shipped.

Many of these features are described in this chapter.

The drive contains a large, segmented buffer in RAM that is used to maximize the throughput for sequential read streams. This cache is organized as a number of separate segments.

The buffer is also used to combine or coalesce I/O data on a track-by-track basis, prior to accessing the disk media. Performance is greatly improved by gathering and combining the data on a track basis rather than on a command basis. With this unique system, multiple commands for a particular track are coalesced in the buffer before accessing the media.

As a result, multiple I/O commands, that address a single track are logically combined to avoid multiple disk accesses on multiple rotations. Since most reads and writes are sequential, this greatly increases throughput for most applications.

6.2 WRITE-BACK CACHING

The drive supports write-back caching. In this case, writes are considered complete when the data is loaded into the cache, before it is written to the media. This allows the host to go on to other tasks, while the drive assumes responsibility for the data and ensuring that it is written to media as soon as possible. If write-back caching is employed, it is recommended that an uninterruptible power source (UPS) be provided for optimal data security.

6.3 ZERO LATENCY READ/WRITE

An average of half a revolution of latency is saved by starting to read or write as soon as a seek settles on the designated track. If reading, the initial data read is stored in a buffer (cache) until the actual starting address data block is reached and read. The starting address data is then transferred to the bus as the rest of the data is transferring from media to the buffer. When the initial data address read is reached again, that data in the buffer is transferred to the bus. This ensures that the track data goes out in media order and eliminates the rotational latency associated with waiting for the starting address to rotate to the head.

Writing works essentially the same way only in reverse. The data is first transferred to the drive's buffer. After seeking to the desired track, a write pointer is set to the current location of the head on the track and writing is commenced from the buffer, starting at the corresponding place in the data.

6.4 DISCONNECT-RECONNECT

System throughput can be improved by disconnecting the drive from the initiator during physical positioning operations, thereby freeing up the SCSI bus for other I/O processes. After the drive has determined that there will be a delay, it disconnects itself from the SCSI bus by sending a DISCONNECT message and enters the BUS FREE phase.

When the drive is ready to resume data transfer, it arbitrates for the SCSI bus and, after winning arbitration, reconnects to the initiator and sends an IDENTIFY message via the MESSAGE IN phase. This revives the I_T_L nexus (initiator-target-logical unit connection) so the initiator can retrieve the correct set of data pointers for the I/O process. The initiator restores the active pointers to their most recent saved values, and the drive continues to finish the original I/O process.

Disconnect-Reconnect is controlled with the Mode Select Command (15h) for the Disconnect/Reconnect (02h) mode page.

6.5 TRACK AND CYLINDER SKEWING

The disk drive improves data throughput by skewing track and cylinder addresses. When the drive switches heads or tracks, or both, to access sequential data, the rotation of the disk media allows one or more physical blocks to pass by the read/write head before the head is ready to continue the transfer. Skewing adjusts the block addresses so that the head switch and settle times and the media rotation coincide to bring the head exactly to the next logical block. This minimizes rotational latency (and increases throughput) when data is accessed sequentially.

6.6 AVERAGE ACCESS TIME

A 10,000 rpm rotation speed yields an average latency of 3 ms.

6.7 EMBEDDED SERVO SYSTEM

Embedded servo information is written in a spoke configuration on every track, on every disk surface. The spokes (or headers) consist of quadrature analog patterns and

digital address data. The digital portions of the spoke data are read and used to locate the desired track, spoke, and head number. The quadrature analog signal portion is detected and used by a servo feedback control loop to precisely position the head on the track center.

6.8 DATA INTEGRITY AND SECURITY

The disk drives use a combination of parity checking, error detection coding (EDC), error correction coding (ECC), and checkpointing to protect stored data from media errors, transfer or addressing errors, or errors introduced during block reallocation.

6.8.1 Media Error Protection

To ensure that data read is the same as data written, the drive computes and appends an Error Correction Code (ECC) to each block of data stored. The drive uses a 352-bit Reed Solomon code with a 4:1 interleave, which can correct up to 20 bytes in each block.

The drive can also correct up to 2 bytes per interleave (up to 8 per block) in hardware (“on-the-fly”), with no loss in throughput.

6.8.2 Transfer Error Protection

An end-to-end error detection code (EDC) protects data from any errors introduced by internal buses, the disk controller chip, the data cache, or the SCSI interface.

An EDC is calculated and added to each data block as the data arrives from the SCSI bus (after SCSI bus parity is checked). The EDC is stored with the data and protected by the block ECC for added security. On reading or writing, the EDC is checked as the data is transferred between buffer RAM and the media or the SCSI bus.

6.8.3 Addressing Error Protection

Each data block on the media is identified and located by a servo spoke address. The spoke address consists of a two-byte word. Each spoke has multiple copies of the least significant bytes of the address. The disk hardware requires that a majority of the copies agree and that the result agrees with the expected head, track, and spoke number, before it will read or write the data.

To further protect against addressing errors, the logical address (LBA) of the data is added to the EDC of each block. If data is written to the wrong block and subsequently read, or read from the wrong block, the error will be flagged.

The hardware does not allow a blind read of a data block; the firmware must request specific data blocks. Even if the head selection hardware malfunctions, it is not possible for the drive to return data from the wrong head.

6.8.4 Data Sector Reallocation Error Protection

In any SCSI disk drive, bad blocks may be reallocated. However, a power failure or unrecoverable data could threaten data integrity during a block reallocation.

The reallocation and defect list storage algorithms prevent a reallocation from being lost due to a power failure. Once a reallocation starts, the information about the data block to be moved is stored on the media. As the reallocation progresses, checkpoint

information about the reallocation (such as the new destination and the data in transit) is recorded. If the reallocation is interrupted, the process can automatically continue from the last checkpoint without data loss.

If a block reallocation is performed on unrecoverable data (rare), a Bad Data Mark is set in the new location. Thus, there is no chance of undetected bad data being generated from a reallocation.

6.8.5 Data Verification

Data can be verified in two ways. First, the integrity of data in selected areas of the media can be checked using the VERIFY command. Secondly, the WRITE AND VERIFY command writes host-supplied data and then verifies that it was written correctly.

The Verify Error Recovery Page of the MODE SELECT COMMAND controls the drive's response to error conditions that arise during the VERIFY command and during the verify operation of the WRITE AND VERIFY command.

6.9 TAGGED COMMAND QUEUING

The drive supports all three SCSI-defined queue tag messages. When queuing is enabled (default condition), the drive can accept any of the following:

- Simple - specifies that the command is to be placed in the drive's command task set (queue). If several commands are present, the drive may reorder them to increase overall throughput.
- Head of Queue - directs the drive to place the command at the beginning of the queue, to be executed next. Consecutive commands with Head of Queue specified are executed in last-in-first-out (LIFO) order. Queue depth is 64.
- Ordered - specifies that commands in the drive's task set are to be executed in the order received (FIFO).

The DQue bit (disable queuing) of the Control Mode page can be used to disable tagged command queuing.

6.10 COMMAND REORDERING

The drive uses Optimized Reordering Command Algorithm (ORCA) and tagged command queuing to provide greater throughput by reordering any commands queued on the drive to minimize rotational latency, seek time and head selection time.

6.11 BANDED RECORDING

In banded recording, the disk is divided into multiple bands (also called partitions, notches, or bit-zoned areas). Starting at the inner band, each band further out has more blocks per track (a higher recording frequency). This use of multiple-frequency recording increases the capacity of the drive.

6.12 **DIAGNOSTICS**

The drive has extensive diagnostic capabilities, including those described below: power-on self-test, periodic self-adjustments, and host diagnostics. Maxtor uses self-diagnostic tests and sense data tracking to manage drive errors, in addition to logical block address (LBA) revectoring.

6.12.1 **Power On Self Test (POST)**

The disk drive performs a self-diagnostic test immediately at power up. Both the Fault LED and a Busy LED briefly illuminate during spin-up as a lamp test.

If the drive passes the self-tests (and the spin-up jumper is installed), it spins up and performs further diagnostics on the Head Disk Assembly (HDA). The drive indicates “not ready” until the HDA tests are completed successfully.

The drive diagnostics test the following:

- RAM and ROM chips
- Buffer memory
- Gate array chips
- Read/write encoding circuits
- Basic servo functions

After a SELECTION TIME following power-on, the disk drive is able to respond with appropriate status and sense data to the TEST UNIT READY, INQUIRY, and REQUEST SENSE commands.

6.12.2 **Periodic Self-Adjustments**

During normal drive operation, certain periodic functions occur automatically so the drive can verify and maintain the integrity of its operations and optimize performance. The operations performed include those listed below.

- Move/Switch idle heads
- Adaptive seek profile calibration
- ROM checksum

Such periodic functions can make the drive appear to be active when no commands are outstanding, because the drive appears to be seeking. Periodic calibrations occur only after the drive has been inactive for 30 seconds.

Note: If a media access command is received during a calibration, the drive aborts the calibration process and acts on the command.

If desired, calibrations can be suspended to reduce drive overhead by setting the Manual Calibration bit in the (Vendor) Special Function Control page. If this is done, manual calibrations can still be initiated by performing a REZERO UNIT command.

6.12.3 Host Diagnostics

The Host Processor can command the drive to perform diagnostic tests and report status using RECEIVE and SEND DIAGNOSTIC RESULTS commands. Certain parameters can be sent to the drive to alter or enhance the performance of the diagnostic tests specified.

6.12.4 Log Sense

Maxtor implements multiple LOG SENSE pages which logically group together parameters related to disk and subsystem activities. Monitoring this data enables users to observe drive performance and possibly predict failures.

6.13 ERROR RECOVERY

The mode pages control all error recovery. Any error recovery that is enabled is attempted without host intervention.

6.13.1 Seek Error Recovery

Seek error recovery is automatic if the hardware is physically able to position to the correct track.

6.13.2 Data Read Errors

Data read errors are recovered, if possible, by one or more of the following methods:

- Correction of the data “on-the-fly” with ECC hardware.
- Re-try the operation.
- Reload read channel registers to calibrated values.
- Off-track Recovery. Use different off-track offsets and re-try.
- Use a fixed gain while retrying.
- Margin servo thresholds to facilitate seek and detent.
- Margin the error tolerance of sync mark detection.
- Force sync mark detection while retrying.
- Skip Sector Rereads. Retries, each without one of eight earlier servo blocks.
- Correction of the bad data with the software ECC algorithm.
- Switching the bias current of the MR head between retries.
- TA baseline compensation, and other TA recovery methods.

6.13.3 Reallocation of Bad Blocks

Most of the spare blocks are located on the inner cylinders of the drive. All defective blocks are in-line spared at the factory or at format time. Some spare blocks are

allocated at the end of each band, and are reserved for grown, revectoring blocks. All grown defects are reallocated to one of the spare blocks, using a “nearest neighbor” algorithm.

If bad block replacement is enabled (via the AWRE and ARRE bits in the Read-Write Error Recovery Page), blocks that are difficult to read or write may be replaced after multiple retries.

If the bad block data is recoverable, the drive first verifies that the original block location is bad with multiple write tests. If the write tests fail, the original data is reassigned to a new location.

6.14 THE ULTRA160 and ULTRA320 LOW VOLTAGE DIFFERENTIAL (LVD) SCSI INTERFACE

SCSI-3 provides increased performance and versatility to SCSI disks. LVD uses lower level voltage swings, differential signaling, and double transition (DT) clocking to allow a maximum bus speed of 160 or 320 MB/second in wide SCSI configurations.

The Maxtor Atlas 10K V disk drive supports the LVD/MSE (Multi-Mode LVD and Single-Ended) interface standard. When installed on a bus with only LVD devices the Maxtor Atlas 10K V disk drive operates in LVD mode at transfer rates up to 160 or 320 MB/second. It is fully compatible with Ultra2 LVD devices operating at slower transfer rates. If an Maxtor Atlas 10K V disk drive is installed on a bus with one or more single-ended (SE) devices, it automatically switches modes to operate in SE mode at the slower, legacy data rates.

Differential SCSI operation has the advantage of higher reliability of data transfer through increased immunity to electrical noise. Differential signaling uses a two-wire active system in which current from each wire flows in opposite directions and reverses direction for a signal transition. The direction of current flow determines whether the voltage levels on the two wires have a high/low or low/high relationship to each other. The common mode, or reference, level for the two LVD signals is +1.25 volts. High and low is defined relative to this level. LVD permits a wide range of signal amplitudes. The Maxtor Atlas 10K V interface drivers are designed to provide a high to low range on each signal of approximately 400 millivolts in a nominal configuration, resulting in a differential signal of approximately 800 millivolts.

Single-ended SCSI, in contrast, uses one-wire active signaling with the signal return wire connected to ground. The active signal range is switched between a high in the range of +2.4 - 3.0 volts and a low in the range of 0 - +.5 volts. LVD/MSE SCSI devices in single-ended mode do not have the signal return wires connected directly to ground, as these same pins must be driven to LVD levels when the device is operating in LVD mode. Instead, when in single-ended mode, the device turns on a driver transistor for each return line that acts as a switch to connect the signal return to ground.

6.14.1 DIFFSENS

The DIFFSENS signal in the SCSI bus defines the current operating mode of the SCSI bus. A level below +.5 volts defines a single-ended bus. A level between +.7 and +1.9 volts defines an LVD bus. A level above +2.4 volts defines a High Voltage Differential (HVD) bus. Before LVD was introduced, HVD was simply called Differential SCSI.

All single-ended-only devices connect the DIFFSENS signal to ground. On an LVD bus, the terminators drive the DIFFSENS signal to +1.3 volts unless some device is holding the signal at ground. HVD devices pull DIFFSENS up to +5 volts via a resistor.

The Maxtor Atlas 10K V disk drive is designed with a circuit that monitors the level of DIFFSENS at power-on. The drive then enables the appropriate circuits for LVD or single-ended operation, or disables its interface completely if it detects an HVD level on DIFFSENS.

The two normal operating environments for the Maxtor Atlas 10K V disk drive are an all-LVD or (LVD/MSE) bus, in which case it operates in its LVD mode, or a bus that contains at least one single-ended device, in which case the Maxtor Atlas 10K V disk drive operates in single-ended mode.

HVD mode is not supported by the Maxtor Atlas 10K V disk drive. Signal levels on an HVD bus can reach as high as +15 volts due to allowed common mode transients.

<p>CAUTION: The LVD and HVD alternatives are mutually exclusive and damage may occur to SCSI devices if an LVD device is plugged into an operating HVD bus.</p>
--

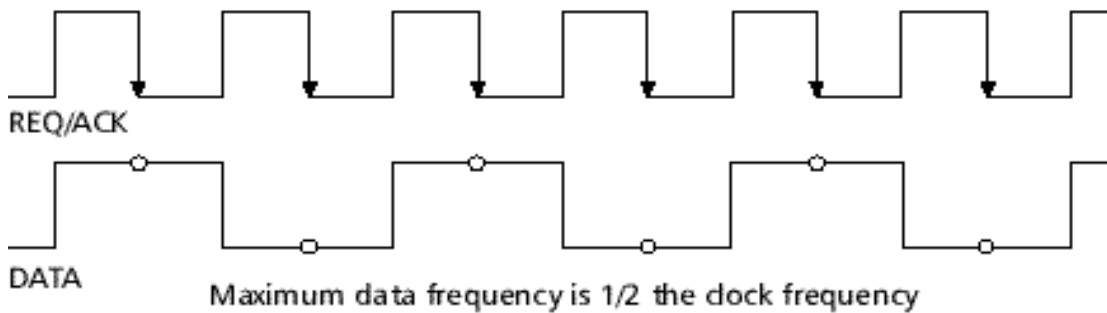
6.14.2 DIFFSENS SWITCHING

After establishing its initial operating mode, the Maxtor Atlas 10K V disk drive continues to monitor the DIFFSENS signal. If a change in level is detected, such as might occur if a device is added to or removed from the bus, the DIFFSENS circuit detects the change and presents a mode change interrupt to the drive firmware. The operating mode of the interface switches to the new mode 100 msec after the change is detected, providing the new level has remained constant. This delay is required to prevent spurious mode switches due to noise on the DIFFSENS line.

6.14.3 DT AND ST CLOCKING

With Ultra160 and Ultra320 SCSI interfaces, Maxtor has introduced a new method of clocking data on the SCSI bus. This method, called Double Transition (DT) clocking, uses both the rising and falling edges of the REQ and ACK signals to clock data into the receiving device. This has the advantage of allowing the REQ and ACK signals to run at half the frequency that would be required under the traditional Single Transition (ST) SCSI clocking scheme for the same data rate. Lower frequency clocks result in a more reliable bus. Figure 6-1 illustrates the difference between DT and ST clocking.

Single Transition (ST) Clocking



Double Transition (DT) Clocking




Figure 6-1 *ST and DT CLocking*

The Maxtor Atlas 10K V disk drive can operate at speeds up to 80 MB/sec in either ST or DT modes, depending on the capabilities of the host adapter. Transfer rates above 80 MB/sec are only supported in DT mode, as required by the SPI-3 SCSI standard.

6.14.4 CYCLIC REDUNDANCY CHECKING

When transferring data with DT clocking on the SCSI bus, a significantly improved form of error checking, cyclic redundancy checking (CRC) is used instead of the traditional parity checking. With CRC checking, the device sending data generates a four-byte CRC character based on the contents of the bytes being transferred. The receiving device also generates a CRC character as it receives the bytes. At the end of a data burst, the sending device transfers its CRC character and the receiving device compares the sender's CRC with its own. If the two CRC characters match, then the data burst transferred without error.

In ST mode, the Maxtor Atlas 10K V uses the traditional parity checking method to ensure error free data transfers.

6.14.5 DOMAIN VALIDATION

A new feature (also known as “physical layer integrity checking”) in Ultra160 and Ultra320 SCSI interfaces is the capability of the initiator and target devices to execute a test sequence of commands before customer operation begins to ensure that the SCSI bus can support the desired transfer rate. If errors are found, the initiator can negotiate different transfer parameters until a working configuration is found.

- Domain validation can detect configuration problems such as:
- Wide devices with a narrow bus segment between them
- DT devices connect through an ST expander
- Broken connections

It may also detect poor quality cables, marginal quality signals, and improper termination.

Domain validation is implemented using standard SCSI commands and new versions of the WRITE BUFFER and READ BUFFER commands. The sequence of commands used for domain validation may vary from system to system. The selection of commands used is entirely under control of host firmware and software.

6.14.6 Adaptive Active Filter (AAF)

Also known as “receiver equalization with filtering,” AAF is a closed-loop method of improving received signal quality by amplifying the fundamental frequency of the signal while filtering noise and other undesirable components. Devices implementing AAF establish the gain of its amplifiers by setting the amplitude of the high-frequency portion of the training pattern to be the same as the low-frequency portion at the beginning of the training pattern. Using the training pattern to perform this adjustment of signal amplitude provides for an inherent closed-loop system that can adjust signal quality for different cable plants and changes in system conditions (e.g., when a new device is added to a system causing the electrical characteristics of the cable plant to change). AAF settings may be adjusted as often as necessary because either the initiator or target may initiate the training pattern sequence. A receiver may disable transmitter pre-compensation in a transmitter as AAF performs better in the configuration.

6.14.7 Transmitter Pre-Compensation with Cutback

This is an open-loop method of compensating for some of the signal loss that is most severe on the first part of a signal's transition.

6.14.8 Free Running Clock (FRC)

FRC is used to improve integrity of the clock signal by removing inter-symbol interference (ISI). By having a clock running at a constant neutralized. The free running clock is restricted for use with DT information unit transfers at 320 megabytes per second.

6.14.9 Skew Compensation

Skew compensation of data signals relative to the clock signal; an Ultra320 SCSI device can establish skew compensation simultaneously for each of the received transitions on the data lines so that they occur at the correct time relative to the clock.

6.14.10 Information Unit Transfers

IUT (or "IU Transfers," also known as "packetized transfers") is a method to encapsulate non-data information (like commands sent from the initiator to the target and status sent from the target to the initiator) into packets and transfer those packets at the maximum negotiated data rate of up to 320 megabytes per second. A method to transfer packets for a number of I/O processes without an intervening physical disconnection. A method to minimize overhead by eliminating several bus phase changes per I/O process.

6.15 HOT PLUGGING/REMOVAL AND INSERTION OF SCSI DEVICES

The Maxtor Atlas 10K V SCA low voltage differential (LVD) disk drives are designed for use in "Hot Swap" applications within a properly designed and configured SCSI system. With care, system integrators can design storage arrays and/or SCSI buses using Maxtor Atlas 10K V disk drives that are "Hot Swap Case 4" tolerant.

Case 4 is defined as follows:

- Devices are powered and the bus may have active I/O processes ongoing, but the device being removed or inserted must be idle (no ongoing I/O processes during the insertion/removal).
- Ground connections to the drive must be made and maintained for 1 millisecond before, during, and 1 millisecond after the insertion/removal. The SCA-2 connector used on the Maxtor Atlas 10K V disk drive meets this requirement.
- Device circuitry connected to bus pins must remain "glitch" free during power up or power down.

Note: In a multimode environment, any insertion or removal that changes the bus mode causes a transceiver mode change reset event.

6.15.1 System Considerations

It is not possible for the suppliers of various system components (e.g., targets, initiators, backplanes, terminators) to guarantee that a system can operate under "Hot Swap Case 4" conditions. Therefore, the system integrator bears the responsibility for ensuring that the system can meet "Hot Swap Case 4" operational criteria.

There are two metrics key to ensuring "Hot Swap Case 4" tolerance:

1. When a drive is being hot inserted into an active bus, a smaller capacitance on the SCSI bus pins create a smaller (that is, a lesser charge) transient spike on the bus.
2. When the drive is the closest drive to the connector at which another drive is being hot inserted, it is desirable that the input receivers ignore short low amplitude transients (usually via a low pass filter).

LVD SCSI devices may require more stringent system design to tolerate transients that occur during Case 4 insertion or removal. System integrators should consider the following when designing their backplanes or buses:

1. Larger connector-to-connector spacing causes the transients associated with hot swapping to be attenuated dramatically before reaching an operating drive.
2. Larger SCSI signal-trace capacitance (that is, lower impedance) is preferable for minimizing transient spikes, but can cause other bus problems unrelated to hot swapping. Draft standard SCSI-3 SCSI SPI-2 (SCSI Parallel Interface), Revision 18, Section 6.6.2 illustrates the trade-offs between trace capacitance and drive-to-drive spacing.
3. The universe of drives needs to be homogeneous. A 15 pF drive hot-inserted into a rack of 15 pF drives presents no difficulty. A 30 pF drive inserted into a rack of 10 pF drives with very close connector spacing, however, is more likely to cause problems.

6.15.2 APPLICABLE SCSI PHYSICAL DOCUMENTS

If additional information is required, refer to the following documentation:

- SCSI-3 SPI (SCSI Parallel Interface), revision 15, now ANSI Standard No. X3.253.1993
 - Device Insertion/Removal classes
- SCSI-3 Fast-20, revision 6, now ANSI Standard No. X3.277.1996
 - Documents 20 Mega-transfer/second operation
 - Defined active negation
- SPI-4 Specification Revision 10 Defines LVD signaling
 - Defines Multimode (LVD/SE) transceivers
 - Defined LVD active terminators
 - Defines Multimode (LVD/SE) terminators
 - Defines new functionality for the DIFFSENS pins
 - Documents 40 Mega-transfer/second operation
 - Documents LVD connector pin assignments for LVD buses
 - This document is a superset of SCSI standards, incorporating contents of both SPI and Fast-20

Appendix A

QUICK REFERENCE

Appendix A is a synopsis of SCSI commands, associated pages, messages, sense keys and status codes used in this manual.

Note: Reserved fields in CDBs and Parameters have a value of 0; Reserved fields in states and other parameters sent to an initiator are set to 0.

A.1 SCSI-2/SCSI-3 Equivalent Terminology

Table A-1 contains the SCSI-3 equivalent terms for common terms used in SCSI-2 documents.

Table A-1 SCSI-2/SCSI-3 Equivalent Terminology

SCSI-2	SCSI-3
abort	abort task set
abort tag	abort task
bus device reset	target reset
clear queue	clear task set
command complete	task complete
continue I/O process	continue task
head of queue tag	head of queue
ordered queue tag	ordered
incorrect initiator connection	overlapped commands
I/O process	task
phase	service
queue	task set
queue full	task set full
simple queue tag	simple

A.2 **SCSI Commands and Messages**

Table A-2 and Table A-3 contains the SCSI-3 commands and messages and their corresponding op codes.

Table A-2 SCSI-3 Quick Reference – Commands

Command	Op Code (Hex)
CHANGE DEFINITION	40
FORMAT UNIT	04
INQUIRY	12
LOG SELECT	4C
LOG SENSE	4D
MODE SELECT (6)	15
MODE SELECT (10)	55
MODE SENSE (6)	1A
MODE SENSE (10)	5A
PERSISTENT RESERVE IN	5E
PERSISTENT RESERVE OUT	5F
READ (6)	08
READ (10)	28
READ BUFFER	3C
READ CAPACITY	25
READ DEFECT DATA (10)	37
READ DEFECT DATA (12)	B7
READ LONG	3E
READ SKIP MASK	E8
REASSIGN BLOCKS	07
REBUILD	81
RECEIVE DIAGNOSTIC RESULTS	1C
REGENERATE	82
RELEASE (6)	17
RELEASE (10)	57
REPORT DEVICE IDENTIFIER	A3
REPORT LUNS	A0
REQUEST SENSE	03
RESERVE (6)	16
RESERVE (10)	56
REZERO UNIT	01
SEEK (6)	0B
SEEK (10)	2B
SEND DIAGNOSTIC	1D
SET DEVICE IDENTIFIER	A4
START/STOP UNIT	1B
SYNCHRONIZE CACHE	35
TEST UNIT READY	00
VERIFY	2F
WRITE (6)	0A
WRITE (10)	2A
WRITE AND VERIFY	2E
WRITE BUFFER	3B
WRITE LONG	3F
WRITE SAME	41
WRITE SKIP MASK	EA

Table A-3 SCSI-3 Quick Reference – Messages

Messages	Msg. Code (Hex)
ABORT	06
ABORT TAG	0D
BUS DEVICE RESET	0C
CLEAR QUEUE	0E
COMMAND COMPLETE	00
CONTINUE I/O PROCESS	12
DISCONNECT	04
EXTENDED MESSAGE	01
HEAD OF QUEUE TAG	21
IDENTIFY	80 – FF
IGNORE WIDE RESIDUE	25
INITIATOR DETECTED ERROR	05
LINKED COMMAND COMPLETE	0A
LINKED COMMAND COMPLETE w/FLAG	0B
MESSAGE PARITY ERROR	09
MESSAGE REJECT	07
NO OPERATION	08
ORDERED QUEUE TAG	22
PARALLEL PROTOCOL REQUEST	04
QUICK ARBITRATION AND SELECTION	55
RELEASE RECOVERY	10
SAVE DATA POINTER	02
SIMPLE QUEUE TAG	20
SYNCHRONOUS DATA TRANSFER REQUEST	01
TARGET TRANSFER DISABLE	13
WIDE DATA TRANSFER REQ.	03

A.3 **SCSI Pages**

Table A-4 lists the SCSI pages, along with the page codes and commands.

Table A-4 SCSI-3 Quick Reference - Pages

Pages	Page Code	Cmd
ASCII Implemented Operating Definition	82	12
ASCII Information	01, 02, 03	12
Buffer Overruns and Underruns	01	4D
Caching	08 *	15, 55
Control Mode	0A *	15, 55
Device Identification	83	12
Disconnect-Reconnect	02	15, 55

Format Device	03	1A
Format Status	08	4D
Implemented Operating Definition	81	12
Last n Error Events	07	4D
Medium Types Supported	0B	1A
Non-Medium Error Counter	06	4D
Notch and Partition	0C	15, 55
Peripheral Device	09	1A, 5A
Port Control Page	19	15, 55
Margin Control Subpage	19 (Sub Page 1)	15, 55
Saved Training Configuration Values Subpage	19 (Sub Page 2)	15, 55
Negotiated Setting Subpage	19 (Sub Page 3)	15, 55
Report Transfer Capabilities Subpage	19 (Sub Page 4)	15, 55
Read Error Counter	03	4D
Read-Write Error Recovery	01	15
Rigid Disk Geometry	04	15, 55
Standard Inquiry Data	**	
Supported Diagnostic Pages	00	1C
Supported Log Pages	00	4D
Supported Vital Product Pages	00	12
Translate Address	40	1C
Unit Attention Control	00	15, 55
Unit Serial Number	80	12
Verify Error Counter	05	4D
Verify Error Recovery	07	15, 55
Write Error Counter	02	4D

Note: * Not supported in SCSI-1/CSS mode

A.4 **Sense Keys**

Table A-5 lists the sense keys.

Table A-5 *SCSI-3 Quick Reference – Sense Keys*

Sense Keys	Sense Code - Hex
NO SENSE	0
RECOVERED ERROR	1
NOT READY	2
MEDIUM ERROR	3
HARDWARE ERROR	4
ILLEGAL REQUEST	5
UNIT ATTENTION	6
DATA PROTECT	7
VENDOR SPECIFIC	9
ABORTED COMMAND	B
MISCOMPARE	E

A.5 **Status Codes**

Table A-6 lists the status codes.

Table A-6 *SCSI-3 Quick Reference – Status Codes*

Status Codes	Status Code - Hex
GOOD	00
CHECK CONDITION	02
BUSY	08
RESERVATION CONFLICT	18
TASK SET FULL	28

Appendix B

NEGOTIATED RATE INFORMATION PAGE REFERENCE

This appendix provides information about the INQUIRY command's Negotiated Rate Information Page that is not provided in Section 5.11.2.7. Specifically, it provides the codes and meanings for the Transfer Period Factor field (byte 4) and a table that explains the allowable settings for the QAS, DT, and IU fields (bits 2, 1, and 0, respectively, of byte 7).

B.1 TRANSFER PERIOD FACTOR FIELD

Table B-1 provides the codes and meanings allowed for the Transfer Period Factor field if the value of PARL (Parallel Protocol, bit 2 of byte 6 of the Negotiated Rate Information Page) is 1.

Table B-2 provides the allowable codes when the value of PARL = 0.

Table B-1 *Transfer Period Factor Field Values When PARL = 1*

Code	Description
00h – 07h	Reserved.
08h	Transfer period = 6.25ns (Fast160 or U320 is latched every 6.25 ns) This code is valid only if the protocol options bits (QAS,DT and IU) of the page have values selected that support double transition (DT) Transfers.
00h – 08h	Transfer Period = 6.25ns (Fast160 or U320 is latched every 6.25 ns) This code is valid only if the protocol options bits (QAS,DT and IU) of the page have values selected that support double transition (DT) Transfers.
09h	Transfer period = 12.5 ns (Fast-80 or U160 is latched every 12.5 ns). This code is valid only if the protocol options bits (QAS, DT, and IU) of the page have values selected that support double transition (DT) data transfers.
0Ah	Transfer period = 25 ns (Fast-40 or Ultra2 data is latched every 25 ns).
0Bh	Transfer period = 30 ns (Fast-40 data is latched every 30 ns).

0Ch	Transfer period = 50 ns (Fast-20 data is latched using a transfer period of less than or equal to 96 ns and greater than or equal to 50 ns).
0Dh – 18h	Transfer period = the period factor x 4 (Fast-20 data is latched using a transfer period of less than or equal to 96 ns and greater than or equal to 50 ns).
19h – 31h	Transfer period = the period factor x 4 (Fast-10 data is latched using a transfer period of less than or equal to 196 ns and greater than or equal to 100 ns).
32h - FFh	Transfer period = the period factor x 4 (Fast-5 data is latched using a transfer period of less than or equal to 1020 ns and greater than or equal to 200 ns).

Table B-2 *Transfer Period Factor Field Values When PARL = 0*

Code	Description
00h – 09h	Reserved.
0Ah	Transfer period = 25 ns (Fast-40 data is latched every 25 ns).
0Bh	Transfer period = 30 ns (Fast-40 data is latched every 30 ns).
0Ch	Transfer period = 50 ns (Fast-20 data is latched using a transfer period of less than or equal to 96 ns and greater than or equal to 50 ns).
0Dh – 18h	Transfer period = the period factor x 4 (Fast-20 data is latched using a transfer period of less than or equal to 96 ns and greater than or equal to 50 ns).
19h – 31h	Transfer period = the period factor x 4 (Fast-10 data is latched using a transfer period of less than or equal to 196 ns and greater than or equal to 100 ns).
32h - FFh	Transfer period = the period factor x 4 (Fast-5 data is latched using a transfer period of less than or equal to 1020 ns and greater than or equal to 200 ns).

B.2 Settings for the QAS, DT, and IU fields

Table C-3 provides valid bit combinations for the QAS, DT, and IU fields of the INQUIRY command's Negotiated Rate Information Page (Section 5.11.2.7 of this manual). Only the combinations shown in the table are valid; all other combinations are reserved.

Table B-3 Settings for the QAS, DT, and IU Fields of the Negotiated Rate Information Page (INQUIRY Command)

QAS	DT	IU	Description
0	0	0	Selects Single Transition phases. Use ST DATA IN and ST DATA OUT phases to transfer data
0	1	0	Use DT DATA IN and DT DATA OUT phases with data group transfers.
0	1	1	Use DT DATA IN and DT DATA OUT phases with information unit transfers.
1	1	1	Use DT DATA IN and DT DATA OUT phases with information unit transfers and use QAS for arbitration.

A

ACCESS – (v) Read, write, or update information stored on a disk or other medium. (n) The operation of reading, writing, or updating stored information.

ACCESS TIME – The interval between the time a request is made by the system and the time the data is available from the drive. Includes the seek time, rotational latency, and command processing overhead time. (See also *seek*, *rotational latency*, and *overhead*.)

ACTIVE I/O PROCESS– An I/O process that is currently in execution (not queued).

ACTUATOR – Also known as the *positioner*. The internal mechanism that moves the read/write head to the proper track. The Maxtor actuator consists of a rotor connected to head mounting arms that position the heads over the desired cylinder. Also known as rotary actuator.

AIRLOCK – A patented Maxtor feature that ensures durable and reliable data storage. Upon removal of power from the drive for any reason, the read/write heads automatically park and lock in a non data area called the landing zone. Airlock allows the drive to withstand high levels of non-operating shock. When power is applied to the drive, airflow created from the spinning disks causes the Airlock arm to swing back and unlock the actuator, allowing the heads to move from the landing zone. Upon power down, the Airlock swings back to the locked position, locking the heads in the landing zone. A park utility is not required to park the heads on drives equipped with Airlock (all Maxtor drives).

ALLOCATION – The process of assigning particular areas of the disk to specific data or instructions. An allocation unit is a group of sectors on the disk reserved for specified information. On hard disks for small computer systems, the allocation unit is usually in the form of a sector, block, or cluster. (See also *allocation unit*.)

ALLOCATION UNIT – An allocation unit, also known as a *cluster*, is a group of sectors on the disk that can be reserved for the use of a particular file.

ANSI – Abbreviation for American National Standards Institute

ASCII – Abbreviation for American Standard Code for Information Interchange

ASIC – Acronym for *Application Specific Integrated Circuit*.

AVERAGE SEEK TIME – The average time it takes for the read/write head to move to a specific location. Calculated by dividing the time it takes to complete a large number of random seeks by the number of seeks performed.

B

BACKUP – A copy of a file, directory, or volume on a separate storage device from the original, for the purpose of retrieval in case the original is accidentally erased, damaged, or destroyed.

BAD BLOCK – A block (usually the size of a sector) that cannot reliably hold data due to a physical flaw or damaged format markings.

BAD TRACK TABLE – A label affixed to the casing of a hard disk drive stating which tracks are flawed and cannot hold data. This list is typed into the low-level formatting program when the drive is installed. Maxtor users can ignore bad track tables since Maxtor's built-in defect-management protections compensate for these flaws automatically.

BANDING – A method of variable frequency recording on the media that attempts to keep the bit density constant, regardless of the track radius. Banding greatly increases capacity and data transfer rates.

BEZEL – A plastic panel that extends the face of a drive so that it covers a computer's drive bay opening. The bezel usually contains a drive-activity LED. Also known as the *faceplate*.

BIT – Abbreviation for binary digit. A binary digit may have one of two values—1 or 0. This contrasts with a decimal digit, which may have a value from 0 to 9. A bit is one of the logic 1 or logic 0 binary settings that make up a byte of data. (See also *byte*.)

BLOCK – In UNIX workstation environments, the smallest contiguous area that can be allocated for the storage of data. UNIX blocks are generally 8 Kbytes (16 sectors) in size. In DOS environments, the block is referred to as a cluster. (Note: This usage of the term block at the operating system level is different from its meaning in relation to the physical configuration of the hard drive. See *sector* for comparison.)

BPI – Bits Per Inch. A measure of how densely information is packed on a storage medium. (See also *FCI*.)

BUFFER – An area of RAM reserved for temporary storage of data that is waiting to be sent to a device that is not yet ready to receive it. The data is usually on its way to or from the hard disk drive or some other peripheral device.

BUS – The part of a chip, circuit board, or interface designed to send and receive data.

BYTE – The basic unit of computer memory, large enough to hold one character of alphanumeric data. Comprised of eight bits. (See also *bit*.)

C

CACHE – Specialized High-speed RAM used to optimize data transfers between system elements with different performance characteristics, e.g., disk to main memory or main memory to CPU.

CAPACITY – The amount of information that can be stored on a hard drive. Also known as storage capacity. (See also *formatted capacity*.)

CLEAN ROOM – An environmentally controlled dust-free assembly or repair facility in which hard disk drives are assembled or can be opened for internal servicing.

CMOS – Acronym for complementary metal oxide semiconductor. A low-power, low-heat, high-density semiconductor.

COMMAND QUEUE – The queue used to store I/O processes.

CONNECT – The initiator function that selects a target to establish a nexus and to start an I/O process. The connection that results is an initial connection.

CONNECTION – An initial connection or reconnection that can occur only between one initiator and one target.

CONTINGENT ALLEGIANCE – A condition generated by a CHECK CONDITION or

COMMAND TERMINATED status during which execution of all tagged I/O processes (for the associated I_T_X nexus) are suspended until the condition is cleared. A target preserves an I_T_X nexus until it is cleared by a hard or soft reset, an abort message or another command for that I_T_X nexus. While the contingent allegiance condition exists, the target maintains sense data for the initiator.

CLUSTER – In DOS environments, the smallest contiguous area that can be allocated for the storage of data. DOS clusters are usually 2 Kbytes (4 sectors) in size.

CONTROLLER – The chip or circuit that translates computer data and commands into a form suitable for use by the hard drive. Also known as disk controller.

CONTROLLER CARD – An adapter containing the control electronics for one or more hard disks. Usually installed in a slot in the computer.

CPU – Central Processing Unit. The microprocessor chip that performs the bulk of data processing in a computer.

CRC – Cyclic Redundancy Check. An error detection procedure that identifies incomplete or faulty data in each sector.

CYLINDER – When disks are placed directly above one another along the shaft, the circular, vertical “slice” consisting of all the tracks located in a particular position.

D

DATA SEPARATOR – The circuit that extracts data from timing information on drives that store a combined data and clock signal.

DEDICATED SERVO – A positioning mechanism using a dedicated surface of the disk that contains timing and positioning information only, as compared to surfaces that are also used for data. (See also *embedded servo*.)

DEFECT MANAGEMENT – A technique ensuring long-term data integrity. Consists of scanning disk drives both at the factory and during regular use, de-allocating defective sectors before purchase and compensating for new defective sectors afterward.

DIFFERENTIAL TERMINATION – One of two methods for terminating the SCSI bus, characterized by a two-line signal: a (+) signal and a (-) signal. Used to improve noise resistance when long bus cables are required. See also single-ended termination.

DISK – In general, any circular-shaped data-storage medium that stores data on the flat surface of the platter. The most common type of disk is the magnetic disk, which stores data as magnetic patterns in a metal or metal-oxide coating. Magnetic disks come in two forms: floppy and hard. Optical recording is a newer disk technology that gives higher capacity storage but at slower access times.

DISK CONTROLLER – A plug-in board, or embedded circuitry on the drive, that passes information to and from the disk. The Maxtor hard disk drives all have controllers embedded on the drive printed-circuit board. (See also *controller*.)

DMA – Direct Memory Access. A process for transferring data directly to and from main memory, without passing through the CPU. DMA improves the speed and efficiency by allowing the system to continue processing even while new data is being retrieved.

DOS – Disk Operating System. The most common operating system used in IBM PCs. Manages all access to data on the disk.

DRIVE – Short form of *disk drive*.

DRIVE GEOMETRY – The functional dimensions of a drive, including the number of heads, cylinders, and sectors per track. (See also *logical format*.)

E

ECC – Error Correction Code. The incorporation of extra parity bits in transmitted data in order to detect errors that can be corrected by the controller.

EDC – Abbreviation for error detection code. An end-to-end cyclic redundancy code that is used to protect data from errors that may be introduced by the connecting busses, the disk controller chip, the data cache, or the SCSI interface in a disk drive. See also, *CRC*.

ELECTRONICS MODULE – The printed circuit board containing the drive electronics.

EMBEDDED SERVO – A timing or location signal placed on tracks that store data. These signals allow the actuator to fine-tune the position of the read/write heads.

ENCODING – The conversion of data into a pattern of On/Off or 1/0 signals prior to being written on the disk surface. (See also *RLL* and *MFM*.)

EPROM – Erasable Programmable Read-Only Memory. An integrated circuit memory chip that can store programs and data in a non-volatile state. These devices can be erased by ultraviolet light and reprogrammed with new data.

EXTENDED CONTINGENT ALLEGIANCE – A condition generated by a target and announced with a *INITIATE RECOVERY* message. The mode is used when multi-step extended error recovery procedures are required to prevent interference as from other initiators in a multi-initiator system. See also *contingent allegiance*.

EXTERNAL DRIVE – A drive mounted in an enclosure separate from the computer system enclosure, with its own power supply and fan, and connected to the system by a cable.

F

FALSE – Refers to the logical-zero (or “negated”) state of a binary number. See also, signal negation.

FAST SCSI – The SCSI protocol that governs fast synchronous transfer. See also Synchronous Transfer.

FCI – Flux Changes per Inch. The number of magnetic field patterns that can be stored on a given area of disk surface, used as a measure of data density. (See also *BPI*.)

FIELD – A logically related group of one or more physically contiguous bits.

FILE SERVER – A computer that provides network stations with controlled access to shareable resources. The network operating system is loaded on the file server, and most shareable devices (disk subsystems, printers) are attached to it. The file server controls system security and monitors station-to-station communications. A dedicated file server can be used only as a file server while it is on the network. A non dedicated file server can be used simultaneously as a file server and a workstation.

FIRMWARE – Permanent instructions and data programmed directly into the circuitry of read-only memory for controlling the operation of the computer. Distinct from software, which can be altered by programmers.

FLUX DENSITY – The number of magnetic field patterns that can be stored in a given length of disk surface. The number is usually stated as flux changes per inch (FCI), with typical values in the thousands. (See also *FCI*.)

FLYING HEIGHT – The distance between the read/write head and the disk surface, made up of a cushion of air that keeps the two objects from touching. Smaller flying heights permit denser data storage but require more precise mechanical designs. Also known as fly height.

FORMAT – To write a magnetic track pattern onto a disk surface, specifying the locations of the tracks and sectors. This information must exist on a disk before it can store data.

FORMATTED CAPACITY – The amount of room left to store data on a disk after writing sector headers, boundary definitions, and timing information during a format operation. The size of a Maxtor drive is always expressed in formatted capacity, accurately reflecting the usable space required.

FORM FACTOR – The industry standard that defines the physical, external dimensions of a particular device. For example, most Maxtor hard disk drives use a 3 1/2-inch form factor.

FULL HEIGHT – Term used to describe a disk that occupies the vertical space (3.250 inches) of a full sized 5.25-inch drive. See also half height, low profile.

G

GIGABYTE (GB) –1 billion bytes. The total accessible capacity varies depending on the operating environment.

Glist – Abbreviation. When the drive encounters a defective sector during regular use, it adds the address of the sector to the Grown Defect List (Glist). See also Plist.

GUIDE RAILS – Plastic strips attached to the sides of a hard disk drive in an IBM PC/AT or compatible computer so that the drive easily slides into place.

H

HALF-HEIGHT – Standard drive size equivalent to half the vertical space of a 5 1/4-inch drive.

HARD DISK – A type of storage medium that retains data as magnetic patterns on a rigid disk, usually made of an iron oxide or alloy over a magnesium or aluminum platter. Because hard disks spin more rapidly than floppy disks, and the head flies closer to the disk, hard disks can transfer data faster and store more in the same volume.

HARD ERROR – A data error that persists when the disk is re-read, usually caused by defects in the physical surface.

HARD-SECTORED – The most common method of indicating the start of each sector on a disk, based on information located in the embedded servo. This method is more precise than soft-sectored techniques and results in lower overhead. (See also *soft-sectored*.)

HEAD – The tiny electromagnetic coil and metal pole used to create and read back magnetic patterns on the disk. Also known as read/write head.

HEAD CRASH – Damage to the read/write head, usually caused by sudden contact with the disk surface. Head crash can also be caused by dust and other particles.

HEAD ASSEMBLY – The tiny electromagnetic or magneto-resistive element used to write and read back the magnetic patterns of data on the recording media. See also Diamond Head.

HEAD DISK ASSEMBLY (HDA) – The assembly made up of the spindle motor, spindle, head positioner (actuator), head preamp electronics, and magnetic disk assembly.

HIGH-CAPACITY DRIVE – By industry conventions typically a drive of 100 megabytes or more.

HIGH-LEVEL FORMATTING – Formatting performed by the operating system to create the root directory, file allocation tables and other basic configurations. (See also *low-level formatting*.)

HOME – Reference track used for recalibration of the actuator. Usually the outermost track (track 0).

HOST ADAPTER – A plug-in board that acts as the interface between a computer system bus and the disk drive.

I

INITIALIZATION – See *low-level formatting*.

INTERFACE – A hardware or software protocol, (contained in the electronics of the disk controller and disk drive) that manages the exchange of data between the drive and computer. The most common interfaces for small computer systems are AT (also known as IDE) and SCSI.

INTERLEAVE – The arrangement of sectors on a track. The Interleave Factor is the number of sectors that pass beneath the read/write heads before the next sector arrives. For example, a 3:1 interleave factor means that the heads read a sector, then let two pass by before reading another, requiring three full revolutions of the disk to access the complete data track. Maxtor drives have an interleave factor of 1:1, allowing the system to access a full track of data in a single revolution.

INTERLEAVE FACTOR – The number of sectors that pass beneath the read/write heads before the next numbered sector arrives. When the interleave factor is 3:1, a sector is read, two pass by, and then the next is read. It would take three revolutions of the disk to access a full track of data. Maxtor drives have an interleave of 1:1, so a full track of data can be accessed within one revolution of the disk, thus offering the highest data throughput possible.

INTERNAL DRIVE – A drive mounted inside one of a computer's drive bays, or a hard disk on a card installed in one of the computer's expansion slots.

I/O PROCESS – An I/O process consists of an initial connection (nexus) followed by zero or more reconnections, all pertaining to a single command or a group of linked commands. An I/O process normally ends with a BUS FREE phase following a COMMAND COMPLETE or a RELEASE RECOVERY message.

J

JUMPER – A tiny box that slips over two pins on a circuit board, connecting the pins electrically. Some board manufacturers use Dual In-Line Package (DIP) switches instead of jumpers.

K

KILOBYTE (K) – A unit of measure consisting of 1,024 (2^{10}) bytes.

L

LANDING ZONE – A non-data area on the disk's inner cylinder where the heads can rest when the power is off.

LATENCY – The time during which the read/write heads wait for the data to rotate into position after the controller starts looking for a particular data track. If a disk rotates at 3,600 rpm, the maximum latency time is 16.4 milliseconds, and the average latency time is 8.2 milliseconds.

LBA – Abbreviation for Logical Block Address. A sequentially-numbered identity of a multi-byte block of user data. It does not refer directly to the physical location of the data in a disk drive.

LOGICAL FORMAT – The logical drive geometry that appears to the computer system BIOS as defined by the drive tables and stored in CMOS.

LOGICAL UNIT – A physical or virtual peripheral device addressable through a target.

LOOK AHEAD – The process of anticipating events in order to speed up computer operations. For example, the system can buffer data into cache RAM by reading blocks in advance, preparing the system for the next data request.

LOW-LEVEL FORMATTING – The process of creating sectors on the disk surface so that the operating system can access the required areas for generating the file structure. Maxtor drives are shipped with the low-level formatting already completed. Also known as *initialization*.

LOW PROFILE – Describes drives built to the

3 1/2-inch form factor, which are only 1 inch high. The standard form factor drives are 1.625 inches high.

LPS – Low Profile Series.

LOW-VOLTAGE DIFFERENTIAL (LVD) – Is a differential bus technology that combines much of the bus length, noise, immunity, and performance benefits to conventional differential SCSI with the power consumption and cost of single-ended SCSI Interfaces. Ultra2 SCSI (LVD) increase bus data rates to 80 Mbytes/sec, provides differential data integrity, extends the SCSI bus cable length to 12 meters, and provides easy system configuration for up to 15 peripherals.

M

MB – See *megabyte*.

MAGNETO-RESISTIVE – A unique characteristic of some metals, which, when combined, yields a variable electrical resistance to varying magnetic fluxes. The resistance varies directly with the strength of the flux, field, not simply to a change in flux as with magneto-inductance.

MEDIA – The magnetic film that is deposited or coated on an aluminum substrate which is very flat and in the shape of a disk. The media is overcoated with a lubricant to prevent damage to the heads or media during head take off and landing. The media is where the data is stored inside the disk in the form of magnetic flux or polarity changes.

MEGABYTE (MB) – A unit of measurement equal to 1,000 kilobytes, or 1,024,000 bytes. (See also *kilobyte*.)

MEGAHERTZ – A measurement of frequency in millions of cycles per second.

MHz – See *megahertz*.

MICROPROCESSOR – The integrated circuit chip that performs the bulk of data processing and controls the operation of all of the parts of the system. A disk drive also contains a microprocessor to handle all of the internal functions of the drive and to support the embedded controller.

MICROSECOND (μs) – One millionth of a second (.000001 sec.).

MILLISECOND (ms) – One thousandth of a second (.001 sec.).

MTBF – Mean Time Between Failure. Reliability rating indicating the failure rate expected of a product expressed in power on hours (POH). Since manufacturers differ in the ways they determine the MTBF, comparisons of products should always take into account the MTBF calculation method.

MTTR – Mean Time To Repair. The average time it takes to repair a drive that has failed for some reason. This only takes into consideration the changing of the major sub-assemblies such as circuit board or sealed housing. Component level repair is not included in this number as this type of repair is not performed in the field.

N

.NANOSECOND (ns) – One billionth of a second (0.000 000 001 second or 10^{-9} second).

NEXUS – A relationship between two devices that begins with the establishment of an initial connection and ends with the completion of an I/O process. A nexus relationship is further defined by the letters I (initiator), T (target), L (logical unit), R (target routine), x (either L or R), Q (queue tag), and x_y (either an x or a Q)

O

OVERHEAD – *Command overhead* refers to the processing time required by the controller, host adapter, or drive prior to the execution of a command. Lower command overhead yields higher drive performance. *Disk overhead* refers to the space required for non-data information such as location and timing. Disk overhead often accounts for about ten percent of drive capacity. Lower disk overhead yields greater disk capacity.

OVERWRITE – To write data on top of existing data, erasing it.

OXIDE – A metal-oxygen compound. Most magnetic coatings are combinations of iron or other metal oxides, and the term has become a general one for the magnetic coating on tape or disk.

P

PARITY – An error checking technique wherein the sum of the bits in a data byte is used to set or reset an extra parity bit, depending on whether the sum is an odd or an even number. A second parity check after the data is transferred to a new location checks that the new parity bit matches the transferred bit. Parity checking can only detect an odd number of bit errors (1,3,5, etc.). Parity is defined as odd or even, depending on which condition (odd or even) is used to set the parity bit.

PARKING ZONE – A position in a non-data area on a disk platter reserved for resting the heads when power is off. Using this area prevents the heads from touching the surface in data areas upon power down, adding to the data integrity and reliability of the disk drive. The parking zone is most often inside the disk's inner-most cylinder.

PARTITION – A portion of a hard disk dedicated to a particular operating system and application and accessed as a single logical volume.

PERFORMANCE – A measure of the speed of the drive during normal operation. Factors affecting performance are seek times, transfer rate and command overhead.

PERIPHERAL – A device added to a system as an enhancement to the basic CPU, such as a disk drive, tape drive or printer.

PHYSICAL FORMAT – The actual physical layout of cylinders, tracks, and sectors on a disk drive.

PHYSICAL SECTOR – See sector, data.

PLATTER – Common term referring to the hard disk.

POH – Power On Hours. The unit of measurement for Mean Time Between Failure as expressed in the number of hours that power is applied to the device regardless of the amount of actual data transfer usage. (See also *MTBF*.)

POSITIONER – See *actuator*.

PREFETCH – The technique of buffering data into cache RAM by reading subsequent blocks in advance to anticipate the next request for data. Prefetch speeds up disk accesses to sequential blocks of data.

PRML – Partial-Response Maximum-Likelihood. A sophisticated data-detection method that improves data throughput and increases areal density. PRML read channels are used in place of peak-detection read channels.

Q

QUEUE – A prioritized list or set of processes. A queue may be a FIFO (first-in, first-out), LIFO (last-in, first -out) or a random-access entity.

QUEUE TAG – The value associated with an I/O process that uniquely identifies it from other queued I/O processes in the logical unit from the same initiator.

QUEUED I/O PROCESS – An I/O process that is in the command queue.

R

RAM – Random Access Memory. An integrated circuit memory chip that allows information to be stored and retrieved by a microprocessor or controller. The information may be stored and retrieved in any order, and all storage locations are equally accessible.

RAM DISK – A “phantom” disk drive created by setting aside a section of RAM as if it were a group of regular sectors. Access to RAM disk data is extremely fast, but is lost when the system is reset or turned off.

READ – The function of detecting, processing and transmitting data that has been recorded (written) on a magnetic disk or some other media.

READ AFTER WRITE – A mode of operation requiring that the system read each sector after data is written, checking that the data read back is the same as the data recorded. This operation lowers system speed but raises data reliability.

READ VERIFY – A data accuracy check performed by having the disk read data to the controller, which then checks for errors but does not pass the data on to the system.

READ/WRITE HEAD – The tiny electromagnetic coil and metal pole piece used to create and read back the magnetic patterns (write or read information) on the disk. Each side of each platter has its own read/write head.

RECONNECT – The act of reviving a nexus to continue an I/O process. An initiator reconnects to a target using the SELECTION and MESSAGE OUT phases after winning arbitration. The target reconnects to the initiator using RESELECTION and MESSAGE IN phases after winning arbitration.

RECONNECTION – The state which results from a reconnect. It exists from the assertion of BSY in the SELECTION or RESELECTION phase and lasts until the next BUS FREE phase.

REMOVABLE DISK – Generally said of disk drives where the disk itself is meant to be removed, and in particular of hard disks using disks mounted in cartridges. Their advantage is that multiple disks can be used to increase the amount of stored material, and that once removed, the disk can be stored away to prevent unauthorized use.

RLL – Abbreviation for Run Length Limited. An advanced method used to encode (or compress) data for storage on a disk. RLL encoding stores almost 50 percent more data than MFM (Modified Frequency Modulation). RLL code improves MFM code by encoding data bits in groups rather than singly. Many encoding schemes are possible. The best achieve optimum values of code rates, clocking, power spectrum, error propagation and encoding complexity. A number pair (e.g., 1,7) is usually stated with RLL references. It refers to the minimum number of 0's (1) and the maximum number of 0's (7) that can separate data 1's stored on the disk.

ROM – Read-Only Memory. Integrated circuit memory chip containing programs that can be accessed and read but can not be modified.

ROTARY ACTUATOR – The rotary actuator replaces the stepper motor used in the past by many hard disk manufacturers. The rotary actuator is perfectly balanced and rotates around a single pivot point. It allows closed-loop feedback positioning of the heads, which is more accurate than stepper motors.

ROTATIONAL LATENCY – The delay between when the controller starts looking for a specific block of data on a track and when that block rotates around to where it can be read by the read/write head. On average, it is half of the time needed for a full rotation (about 8 ms.).

S

SCA – Single Connector Attachment. A drive connector that combines the signal and power connectors into one. Typically, SCA-type connectors are used in RAID (Redundant Array of Inexpensive Devices) systems.

SCA-2 – An improved Single Connector Attachment that incorporates two electrostatic discharge (ESD) pins on either end of the connector to improve safety when making a connection.

SCSI – Small Computer System Interface. An interface designed for Apple Macintosh systems and UNIX workstations.

SCSI-1 – The first small computer system interface standard.

SCSI-2 – An enhanced small computer system interface as defined and published by ANSI in standard X3.131-1994. One of the major features of SCSI-2 is integrating the controller into the peripheral device, making the device independent of the initiator.

SCSI-3 – A new set of related ANSI standards are currently under development by the X3T10 committee. The old SCSI-2 standard has been divided into a new set of standards for SCSI-3. Some of these [working draft] standards are:

SAM (SCSI Architecture Model) for the architecture;

SBC (SCSI Block Commands) for the disk drive specific command set.

SPC (SCSI Primary Commands) for the primary command set; and

SIP (SCSI Interlocked Protocol) for the protocol layer;

SPI (SCSI Parallel Interface) for the physical layer;

Layering the standards documents in this manner allows substitution of parts of the structure as new technology arrives.

SCSI ADDRESS – The unique address assigned to a SCSI bus device. Normally, the address is assigned and set during system installation with priority being a consideration. A SCSI address is a value of 0 to 7 (or 0 to 15 for “wide” SCSI).

SCSI BUS – A multi-signal parallel bus that interconnects SCSI devices in a daisy-chain fashion.

SCSI DEVICE – A host adapter or a target controller that can be attached to the SCSI bus.

SCSI ID – The bit-significant representation of the SCSI address referring to one of eight (or one of sixteen) data lines.

SECTOR – On a PC hard drive, the minimum segment of track length that can be assigned to store information. On Macintosh and UNIX drives, sectors are usually grouped into blocks or logical blocks that function as the smallest data unit permitted. Since these blocks are often defined as a single sector the terms block and sector are sometimes used interchangeably in this context. (Note: The usage of the term block in connection with the physical configuration of the disk is different from its meaning at the system level. See also *block* and *cluster* for comparison.)

SEEK – A movement of the disk read/write head to a specific data track.

SEGMENTED CACHE BUFFER – A cache buffer that is organized into multiple track lines. Segmenting the cache allows track data to be saved in separate segments when the head is switched rather than having to erase the entire cache. Segmenting the cache enables command reordering on a more efficient track-basis rather than on a command-basis.

SENSE DATA – Information returned to an initiator regarding error conditions in the drive.

SERVO SECTOR TIME – The time (in μsec) between reading one servo header (or “spoke”) and reading the next header.

SERVO SPOKE DATA – In embedded servo technology, digital and quadrature analog servo data is written (embedded) in multiple headers (or “spokes”) on each track on the data-recording area of each disk surface. The digital portion of the spoke data are read and used to locate the correct track, spoke, and head number. The quadrature analog signal portion is detected and used by a servo feedback control loop to precisely position the head on the track

center.

SERVO DATA – Magnetic markings written on the media that guide the read/write heads to the proper position.

SERVO SURFACE – A separate surface containing only positioning and disk timing information but no data.

SETTLE TIME – The interval between the arrival of the read/write head at a specific track, and the lessening of the residual movement to a level sufficient for reliable reading or writing.

SHOCK RATING – A rating, expressed in “G’s”, of how much shock a disk drive can sustain without damage.

SIGNAL ASSERTION – The act of driving a signal to the true (logical-one) state. An asserted signal on the SCSI bus is a low voltage resulting from driving a transistor ON.

SIGNAL NEGATION – The act of driving a signal to the false (0) state, or allowing the cable terminators to bias the signal to the false state by placing the data driver in the high impedance (off) condition. A negated signal on the SCSI bus is a high voltage resulting from all OR-tied transistors being OFF.

SIGNAL RELEASE – The act of allowing the cable terminators to bias the signal to the false state by placing the bus driver in the high impedance (off) condition.

SINGLE-ENDED TERMINATION – One of two methods for terminating the SCSI bus, characterized by signal values that conform to the ANSI-defined levels for SCSI implementation. Typically used for short cable runs. See also differential termination.

SKEW – To shift sector addresses when switching tracks or cylinders to create a slight overlap under the read/write heads. The overlap works to delay the arrival of the next (in sequence) sector from arriving under the selected head until the track or cylinder switch has been done. Skewing minimizes latency time (increasing data throughput) when data on the drive is accessed sequentially.

S.M.A.R.T. – Self-Monitoring and Analysis Reporting Technology. The continual monitoring of the drive's operating condition with reports sent to the host computer. Similar to Maxtor's EWS (Early Warning System) used in some drives.

SOFT ERROR – A faulty data reading that does not recur if the same data is reread from the disk, or corrected by ECC. Usually caused by power fluctuations or noise spikes.

SOFT-SECTORED – Old time-based method of indicating the start of each sector on a disk. Soft-sectored drives require that location instructions be located in the data fields. (See also *hard-sectored*.)

SPINDLE – The drive's center shaft, on which the hard disks are mounted. A synchronized spindle is a shaft that allows two disks to spin simultaneously as a mirror image of each other, permitting redundant storage of data.

SPUTTER – A special method of coating the disk that results in a hard, smooth surface capable of storing data at a high density. Maxtor disk drives use sputtered thin film disks.

STEPPER – A type of motor that moves in discrete steps with each electrical pulse. Stepper were originally the most common type of actuator engine, since they can be geared to advance a read/write head one track per step. However, they are not as fast, reliable, or durable as the voice coil actuators found in Maxtor disk drives. (See also *voice coil*.)

SUBSTRATE – The material underneath the magnetic coating of a disk. Common substrates include aluminum or magnesium alloys for hard drives, glass, for optical disks, and mylar for floppy disks.

SURFACE – The top or bottom side of a disk, which is coated with the magnetic material for recording data. On some drives one surface may be reserved for positioning information.

SYNCHRONOUS DATA TRANSFER – Data transfer that allows transfer of multiple blocks of data between request or acknowledgment signals, since request and acknowledgment signals need not be sent

before and after each block. This method is faster than asynchronous data transfer.

T

TARGET – An SCSI device that performs an operation requested by an initiator.

TARGET ROUTINE – An I/O process directed to a target and not to a logical unit.

THIN FILM – A type of coating allowing very thin layers of magnetic material, used on hard disks and read/write heads. Hard disks with thin film surfaces can store greater amounts of data.

THIRD-STROKE SEEK – A method of evaluating the power requirements of a hard drive. It involves monitoring drive current while seeking to a selected track, waiting for 1/2 a revolution, then seeking to a track that is 1/3 of the total actuator stroke from the first. The seek/wait cycle is usually repeated three times.

TPI – Tracks Per Inch. The number of tracks written within each inch of disk's surface, used as a measure of how closely the tracks are packed on a disk surface. Also known as *track density*.

TRACK – One of the many concentric magnetic circle patterns written on a disk surface as a guide for storing and reading data. Also known as *channel*.

TRACK DENSITY – How closely the tracks are packed on a disk surface. The number is specified as tracks per inch (TPI).

TRACK-TO-TRACK SEEK TIME – The time required for the read/write heads to move to an adjacent track.

TRANSFER RATE – The rate at which the disk sends and receives data from the controller. The sustained transfer rate includes the time required for system processing, head switches and seeks, and accurately reflects the drive's true performance. The burst mode transfer rate is a much higher figure that refers only to the movement of data directly into RAM.

TRUE – Refers to the logical-one (or

“asserted”) state of a binary number. See also, signal assertion.

U

UNFORMATTED CAPACITY – The total number of bytes of data that can be put on a disk drive. Formatting the disk requires approximately 20% of this space to record servo data, addresses, boundary definitions, and timing information. After formatting, user data can be stored on the remaining disk space, known as the formatted capacity.

ULTRA SCSI– Doubles the bandwidth of SCSI Fast. It provides 8-bit (SCSI NARROW) data rates of 20 Mbytes per second and 16-bit (SCSI Wide) data rates of 40 Mbytes per second. Shorter cables may be required. Also known as Fast-20.

UNFORMATTED CAPACITY – The total number of usable bytes on a disk, including the space that will be required to later to record location, boundary definitions, and timing information. (See *formatted capacity* for comparison.)

V

VOICE COIL – A fast and reliable actuator motor that works like a loud speaker, with the force of a magnetic coil causing a proportionate movement of the head. Voice coil actuators are more durable than their stepper counterparts, since fewer parts are subject to daily stress and wear. Voice coil technology is used in all Maxtor drives.

W

Wide SCSI – SCSI protocol permits wide data transfers of 16 or 32 bits.

Wide data transfer – Data transfer that allows information transfer across 16 (or future 32) bits. Faster than standard 8-bit data transfer.

WEDGE SERVO – The position on every track that contains data used by the closed loop positioning control. This information is used to fine tune the position of the read/write heads exactly over the track center.

WINCHESTER DISKS – Former code name for an early IBM hard disk model, sometimes still used to refer to hard drives in general.

WRITE – The operation of recording or storing data on a magnetic or other media. See also, read.

WRITE ONCE – An optical disk technology that allows the drive to store and read back data, but prevents the drive from erasing information once it has been written.

WRITE SAME - A command requests that the drive write the single sector of data transferred by the initiator to the medium multiple times.

INDEX

Numerics

4-pin DC power connector 3-21
512-byte sectors available 4-3

A

abbreviations 1-2
acoustics 4-7
Active Notch 5-51
actuator 4-6
Adapter, 50-pin to 68-pin 2-4
Additional Sense Codes 5-109, 5-113
Additional Sense Length 5-110
Allocation Length 5-11, 5-34, 5-101, 5-107, 5-108, 5-109
altitude2 4-8
ASCII Implemented Operating Definition 5-26
ASCII Implemented Operating Definition Page 5-26
audience definition 1-1
auto head-park method 4-2
average write 4-2
AWRE 5-15, 5-43, 5-47

B

Bit Pointer 5-122
Block Descriptor 5-11, 5-15, 5-40, 5-41, 5-63, 5-64, 5-65, 5-70
Block Length 5-41, 5-110
Buffer Empty Ratio 5-44
Buffer Full Ratio 5-44
Buffer ID 5-92, 5-140, 5-141
Buffer Offset 5-92, 5-140, 5-141
buffer size 4-2
Bus Inactivity Limit 5-44
bus interface connector 3-16, 3-17, 3-20
Byte Transfer Length 5-11, 5-99, 5-142
bytes per sector 4-1

C

Caching 5-46, 5-47, 5-138

Caching Page 5-46, 5-138
CHANGE DEFINITION Command 5-14
Changeable Pages 5-41
clearance 3-24
Clearing Sense Data 5-108
Command Queuing 5-23
component life 4-10
Connect Time Limit 5-44
connector 3-15, 3-16, 3-17, 3-20
contact start/stop cycles 4-2
Contingent Allegiance Condition 5-8
Control Mode 5-8, 5-48, 5-49
Control Mode Page 5-48, 5-49
Correction Span 5-43, 5-46
Cylinder Number of Defect 5-19
Cylinder Skew Factor 5-66

D

data transfer rate 2-2, 4-4
Data transfer Rates 4-2
Defect Descriptors 5-11, 5-12, 5-100
Defect List Format 5-95, 5-96, 5-97, 5-98
Defect List Header 5-11, 5-12, 5-17, 5-18, 5-95, 5-96, 5-97, 5-98, 5-100
Defect List Length 5-17, 5-18, 5-96, 5-98, 5-100
Defect Lists 5-15
Defective Sector Number 5-19
Disconnect Time Limit 5-44
Disconnect-Reconnect Page 5-44
disk errors 4-3, 4-11
disks 4-1
DOS 4-3
Drive Step Rate 5-68
drop 4-10

E

ECC 5-12, 5-15, 5-43, 5-46, 5-99, 5-142
encoding method 4-2
Ending Boundary 5-51
environmental conditions 4-8

error correction method 4-2
Error Reporting 5-7
error type 4-11
ESD 3-4
Extent List Length 5-124

F

faceplate 3-3
Field Pointer Bytes 5-122
Firmware Revision 5-27
Firmware Revision Page 5-27
Flag Bit 5-6, 5-13
Format Device 5-65, 5-66
Format Device Page 5-65
FORMAT UNIT Command 5-15, 5-16,
5-17, 5-21
formatted capacity 4-1, 4-3
full stroke 4-4

G

general 2-1
GList 5-95, 5-96, 5-97, 5-98
Grown Defect List 5-15, 5-17, 5-18, 5-
43, 5-95, 5-96, 5-97, 5-98, 5-100

H

hard reset time 4-4
HARDWARE ERROR 5-112
hardware options 3-6
hardware requirements 2-4
Head Offset Count 5-43
host system 3-7
hot plugging 6-11
humidity 4-8

I

idle 4-6
Immed 5-15, 5-18, 5-133, 5-134
Implemented Operating Definition 5-25,
5-26
Implemented Operating Definition Page

5-25, 5-26
Initialization Pattern Descriptor 5-19
input power connections 3-21
inside zone 4-1
interface connector 3-15
Interleave 5-66

J

J1 12-pin Auxiliary Connector 3-22
Jumper Configurations and Connections
(16-bit Drives) 3-6
jumpers 3-7

K

key features 2-1

L

Landing Zone Cylinder 5-68
Link Bit 5-6, 5-13
Log Pages 5-33
LOG SELECT Command 5-32
LOG SENSE Command 5-33, 5-34
Logical Block Address 5-40, 5-51, 5-89,
5-90, 5-93, 5-99, 5-100, 5-110, 5-
112, 5-127, 5-128, 5-129, 5-134,
5-136, 5-137, 5-138, 5-139, 5-142
low 4-3
low-level format 4-3

M

manual organization 1-1
mating connectors 3-21
Maximum Burst Size 5-44
maximum effective areal density 4-2
maximum linear density 4-1
maximum number of errors 4-11
maximum screw torque 3-24
maximum specifications 4-4
MEDIUM ERROR 5-18, 5-100, 5-111,
5-112, 5-122
Medium Rotation Rate 5-68

INDEX

Medium Type 5-63, 5-65, 5-70
Mode Data Length 5-65
Mode Page Types 5-39
Mode Parameter Block Descriptor 5-40, 5-41
Mode Parameter Header 5-11, 5-39, 5-40, 5-62, 5-64, 5-69, 5-70
Mode Parameter List 5-39
mode select 1-3
MODE SELECT (10) 5-1, 5-62
MODE SELECT (6) Command 5-37, 5-60
MODE SENSE (10) 5-1, 5-69
Module Serial Number 5-24
mounting 3-23, 3-24
multiple drive shipments 3-5

N

naming conventions 1-3
New Operating Definition 5-14
NO SENSE 5-8, 5-111
NOT READY 5-110, 5-111, 5-121, 5-122
Notch and Partition 5-38, 5-50, 5-51
Notch and Partition Page 5-50
Number of Blocks 5-41, 5-134
Number of Cylinders 5-68
Number of Heads 5-68

O

orientation 3-23
outside zone 4-1

P

packing assembly 3-5
packing materials 3-4
Page Length 5-35
Pages Notched 5-51
Parameter Code Reset 5-32
Parameter List 5-11, 5-12, 5-17, 5-18, 5-37, 5-39, 5-130, 5-141
Parameter List Length 5-37, 5-130, 5-141

Parameter Pointer 5-34
Parameter Pointer Control 5-34
PCBA 3-24
performance 2-2
Peripheral Device 5-23
Peripheral Qualifier 5-23
physical shock 3-4
pin assignments 3-22
PList 5-95, 5-96, 5-97, 5-98
PM 4-10
POWER 4-5
power down 4-5
power on 4-4
power requirements 4-6
power sequencing 4-5
Primary Defect List 5-15, 5-17, 5-18, 5-95, 5-96, 5-97, 5-98
printed circuit board 3-7
product overview 2-1
Progress Indication 5-108, 5-121, 5-122
Progress Indication Bytes 5-108

Q

Queue Algorithm Modifier 5-49

R

R/W heads 4-1
READ (10) Command 5-90
READ (6) Command 5-71, 5-89
READ BUFFER Command 5-91, 5-92
READ CAPACITY Command 5-93, 5-94, 5-95, 5-97
READ DEFECT DATA Command 5-95
READ LONG Command 5-99
Read Retry Count 5-43
Read-Write Error Recovery 5-38, 5-42, 5-47
Read-Write Error Recovery Page 5-42
REASSIGN BLOCKS Command 5-100
recording technology 4-1
RECOVERED ERROR 5-43, 5-46, 5-111, 5-112
Recovery Time Limit 5-43, 5-46

Reed-Solomon ECC 2-3
references 1-3
Relative Addressing 5-23, 5-90, 5-93
RELEASE Command 5-104, 5-105
reliability 2-3, 4-10
Remote Busy and Fault Displays (16-bit
Drives, except SCA) 3-9
REQUEST SENSE Command 5-106, 5-
107, 5-108, 5-132
Reservation Identification 5-104, 5-105,
5-124
RESERVE Command 5-124, 5-125
reset 4-5
Response Data Format 5-23
Retry Count 5-122
Returned Logical Block Address 5-93
REZERO UNIT Command 5-127
Rigid Disk Geometry 5-67, 5-68
Rigid Disk Geometry Page 5-67
Rotational Offset 5-68
rotational speed 4-1
RPL Status Change 5-110

S

SCSI configurations 2-1
SCSI ID bits 3-8
SCSI interface configurations 3-3
SCSI-2 specification 1-3
SEEK (10) Command 5-129
SEEK (6) Command 5-128
Segment Number 5-110
SEND DIAGNOSTIC Command 5-130
Sense Data Availability 5-108
Sense Key 5-20, 5-110, 5-111, 5-112, 5-
121, 5-122
Sense Key Information Field Contents 5-
112
Sense Keys 5-111
sequential switch times 4-4
sequential throughput 4-2
shipping container 3-4
shock 4-10
Shock Feet 3-4
space requirements 3-3
specifications 4-1, 4-10

spindle startup 4-6
Standard Inquiry Data 5-22, 5-23
Standard Inquiry Data Page 5-22
start/stop 4-10
Starting Boundary 5-51
Status Reporting 5-109
Status/Error Reporting 5-7
Supplied Format 5-103
supply voltage 4-5
Supported Additional Sense Codes 5-109,
5-113
Supported Diagnostic Page List 5-130, 5-
131
Supported Diagnostics Page 5-102
Supported Diagnostics Pages 5-102
Supported Sense Keys 5-111
SURF 5-66
SYNCHRONIZE CACHE Command 5-
134
Synchronous Data Transfer 5-23

T

TARGET TRANSFER DISABLE 5-23
temperature 4-8
terminology 1-2
TEST UNIT READY Command 5-135
Third Party Device ID 5-104, 5-105, 5-
124, 5-125
timing specifications 4-4
total tracks 4-1
total user sectors 4-1
Track Skew Factor 5-66
tracks density 4-1
Transfer Length 5-11, 5-47, 5-89, 5-90,
5-99, 5-137, 5-138, 5-139, 5-142
transferred data 4-4
transient voltages 4-5
Translate Address 5-101, 5-102, 5-103,
5-130, 5-131
Translate Address Page 5-101, 5-102, 5-
131
Translate Format 5-103
typical specifications 4-4
typographical conventions 1-3

U

Unit Attention Condition Page 5-59, 5-60
Unit Serial Number 5-24
Unit Serial Number Page 5-24
UNIX 4-3
unrecoverable error rate 4-2

V

ventilation 3-23, 3-25
ventilation requirements 3-23
Verification Length 5-136
VERIFY Command 5-136, 5-139
Verify Correction Span 5-46
Verify Error Recovery 5-45, 5-46
Verify Error Recovery Page 5-45
Verify Recovery Time Limit 5-46
Verify Retry Count 5-46
versatility 2-3
vibration 4-10
Vital Product Data 5-11, 5-21, 5-24, 5-28
Vital Product Data Pages 5-24
voltage returns 3-21
Volume Serial Number 5-24, 5-25

W

WRITE (10) Command 5-138
WRITE (6) Command 5-137
WRITE AND VERIFY Command 5-139
WRITE BUFFER Command 5-140, 5-141
WRITE LONG Command 5-142
Write Retry Count 5-43
WRITE SAME 5-12,5-13

Z

zones per surface 4-1